ON UNIFORM CANONICAL BASES IN L_p LATTICES AND OTHER METRIC STRUCTURES

ITAÏ BEN YAACOV

Abstract. We discuss the notion of *uniform canonical bases*, both in an abstract manner and specifically for the theory of atomless L_p lattices. We also discuss the connection between the definability of the set of uniform canonical bases and the existence of the theory of beautiful pairs (i.e., with the finite cover property), and prove in particular that the set of uniform canonical bases is definable in algebraically closed metric valued fields.

Introduction

In stability theory, the *canonical base* of a type is a minimal set of parameters required to define the type, and as such it generalises notions such as the field of definition of a variety in algebraic geometry. Just like the field of definition, the canonical base is usually considered as a set, a point of view which renders it a relatively "coarse" invariant of the type. We may ask, for example, whether a type is definable over a given set (i.e., whether the set contains the canonical base), or whether the canonical base, as a set, is equal to some other set. However, canonical bases, viewed as sets, cannot by any means classify types over a given model of the theory, and they may very well by equal for two distinct types. The finer notion of *uniform* canonical bases, namely, of canonical bases from which the types can be recovered uniformly, is a fairly natural one, and has appeared implicitly in the literature in several contexts (e.g., from the author's point of view, in a joint work with Berenstein and Henson [BBHa], where convergence of uniform canonical bases is discussed).

Definitions regarding uniform canonical bases and a few relatively easy properties are given in Section 1. In particular we observe that every stable theory admits uniform canonical bases *in some imaginary sorts*, so the space of all types can be naturally identified with a type-definable set. We then turn to discuss the following two questions.

The first question is whether, for one concrete theory or another, there exist *mathematically natural* uniform canonical bases, namely, uniform canonical bases consisting of objects with a clear mathematical meaning. A positive answer may convey additional insight into the structure of the space of types as a type-definable set. This is in contrast with the canonical parameters for the definitions, whose meaning is essentially tautological and can therefore convey no further insight. The case of Hilbert spaces is quite easy, and merely serves as a particularly accessible example. The case of atomless probability spaces (i.e., probability algebras, or spaces of random variables), treated in Section 2, is not much more difficult. Most of the work is spent in Section 3 where we construct uniform canonical bases for atomless L_p lattices in the form of "partial conditional expectations" $\mathbf{E}_t[\cdot|E]$ and $\mathbf{E}_{[s,t]}[\cdot|E]$ (defined there). To a large extent, it is this last observation which prompted the writing of the present paper.

2000 *Mathematics Subject Classification.* 03C45, 46B42.

Key words and phrases. stable theory, uniform canonical base, L_p Banach lattice, beautiful pairs.

Author supported by ANR chaire d'excellence junior THEMODMET (ANR-06-CEXC-007) and by the Institut Universitaire de France.

Revision 1143 *of* 28th August 2010.

The second question, discussed in Section 4, is whether the (type-definable) set of uniform canonical bases is in fact definable. We characterise this situation in terms of the existence of a theory of beautiful pairs. In Section 5 we use earlier results to show that for the theory of algebraically closed metric valued fields, the theory of beautiful pairs does indeed exist, and therefore that the sets of uniform canonical bases (which we do not describe explicitly) are definable.

For stability in the context of classical logic we refer the reader to Pillay [Pil96]. Stability in the context of continuous logic, as well as the logic itself, are introduced in [BU].

1. Uniform canonical bases

In classical logic, stable theories are characterised by the property that for every model $\mathcal{M}$, every type $p(\bar{x}) \in \mathrm{S}_{\bar{x}}(M)$ is definable, i.e., that for each formula $\varphi(\bar{x}, \bar{y})$ (say without parameters, this does not really matter) there exists a formula $\psi(\bar{y})$ (with parameters in M) such that for all $\bar{b} \in M$:

$$\varphi(\bar{x}, \bar{b}) \in p \qquad \Longleftrightarrow \qquad \vDash \psi(\bar{b}).$$

In this case we say that ψ is the φ-definition of p, and write

$$\psi(\bar{y}) \quad = \quad d_{p(\bar{x})}\varphi(\bar{x}, \bar{y}).$$

Obviously, there may exist more than one way of writing a φ-definition for p, but since any two such definitions are over $\mathcal{M}$ and equivalent there, they are also equivalent in every elementary extension of $\mathcal{M}$, and thus have inter-definable canonical parameters. In other words, the canonical parameter of the φ-definition of p is well-defined, up to inter-definability, denoted $\mathrm{Cb}_\varphi(p)$. The collection of all such canonical parameters, as $\varphi(\bar{x}, \bar{y})$ varies (and so does $\bar{y}$) is called the *canonical base* of p, denoted $\mathrm{Cb}(p)$. This is, up to inter-definability, the (unique) smallest set over which p is definable. The same holds for continuous logic with some minor necessary changes, namely that the φ-definition may be a definable predicate (i.e. a uniform limit of formulae, rather than a formula), and it defines p in the sense that

$$\varphi(\bar{x}, \bar{b})^p \quad = \quad d_{p(\bar{x})}\varphi(\bar{x}, \bar{b}).$$

We shall hereafter refer to definable predicates as formulae as well, since for our purposes the distinction serves no useful end.

Since canonical parameters are, *a priori*, imaginary elements, the canonical base is a subset of M^{eq}. For most purposes of abstract model theory this is of no hindrance, but when dealing with a specific theory with a natural "home sort", it is interesting (and common) to ask whether types admit canonical bases which are subsets of the model. This is true, of course, in any stable theory which eliminates imaginaries. In continuous logic, this is trivially true for Hilbert spaces, it is proved for probability algebras in [Ben06], and for L_p Banach lattices in [BBHb] (so all of these theories have, in particular, weak elimination of imaginaries, even though not full elimination of imaginaries).

A somewhat less commonly asked question is the following. Can we find, for each formula $\varphi(\bar{x}, \bar{y})$, a formula $d\varphi(\bar{y}, Z)$, where Z is some infinite tuple of variables of which only finitely (or countably) many actually appear in $d\varphi$, such that for every model $\mathcal{M}$, and every type $p(\bar{x}) \in \mathrm{S}_{\bar{x}}(M)$,

$$d_{p(\bar{x})}\varphi(\bar{x}, \bar{y}) \quad = \quad d\varphi\big(\bar{y}, \mathrm{Cb}(p)\big).$$

The scarcity of references to this question is actually hardly surprising, since, first, the question as stated makes no sense, and, second, the answer is positive for every stable theory. Indeed, if we consider $\mathrm{Cb}(p)$ to be merely a set which is only known up to inter-definability, as is the common practice, then the expression $d\varphi\big(\bar{y}, \mathrm{Cb}(p)\big)$ is meaningless. We remedy this in the following manner:

Definition 1.1. Let T be a stable theory. A *uniform definition of types* in the sort of $\bar{x}$ consists of a family of formulae $\{d\varphi(\bar{y}, Z)\}_{\varphi(\bar{x},\bar{y})\in\mathcal{L}}$, where Z is a possibly infinite tuple, such that for each type $p(\bar{x})$ over a model $\mathcal{M} \vDash T$ there exists a tuple $A \subseteq M^{eq}$ in the sort of Z such that for each $\varphi(\bar{x}, \bar{y})$:

$$d_{p(\bar{x})}\varphi(\bar{x}, \bar{y}) \quad = \quad d\varphi(\bar{y}, A).$$

If, in addition, this determines the tuple A uniquely for each p then we write $A = \mathrm{Cb}(p)$ and say that the map $p \mapsto \mathrm{Cb}(p)$ is a *uniform canonical base map*, or that the canonical bases $\mathrm{Cb}(p)$ are *uniform* (in p).

To complement the definition, a (non uniform) *canonical base map* is any map Cb which associates to a type p over a model some tuple $\mathrm{Cb}(p)$ which enumerates a canonical base for p.

First of all, we observe that every uniform canonical base map is in particular a canonical base map. Second, any uniform definition of types gives rise naturally to a uniform canonical base map. Indeed, for each φ we let w_φ be a variable in the sort of canonical parameters for $d\varphi(\bar{y}, Z)$, and let $d\varphi'(\bar{y}, w_\varphi)$ be the corresponding formula. For a type p, let A be a parameter for the original definition, and for each φ let b_φ be the canonical parameter of $d\varphi(\bar{y}, A)$, so $d\varphi(\bar{y}, A) = d\varphi'(\bar{y}, b_\varphi)$. Now let W be the tuple consisting of all such w_φ, so we may re-write $d\varphi'(\bar{y}, w_\varphi)$ as $d\varphi'(\bar{y}, W)$, and let B be the tuple consisting of all such b_φ. Then $d_{p(\bar{x})}\varphi(\bar{x}, \bar{y}) = d\varphi(\bar{y}, A) = d\varphi'(\bar{y}, B)$ for all φ, and in addition this determines B uniquely. Thus $\mathrm{Cb}(p) = B$ is a uniform canonical base map.

Lemma 1.2. *Every stable theory admits uniform definitions of types and thus uniform canonical base maps (in every sort).*

Proof. This is shown for classical logic in, say, [Pil96], and for continuous logic (which encompasses classical logic as a special case) in [BU]. $\blacksquare_{1.2}$

Lemma 1.3. *The image* $\operatorname{img}\mathrm{Cb}$ *of a uniform canonical base map is a type-definable set.*

Proof. All we need to say is that the tuple of parameters does indeed define a (finitely, or, in the continuous case, approximately finitely) consistent type, which is indeed a type-definable property. $\blacksquare_{1.3}$

Lemma 1.4. *Let* Cb *be a uniform canonical base map in the sort $\bar{x}$, and let f be definable function (without parameters) defined on* $\operatorname{img}\mathrm{Cb}$*, into some other possibly infinite sort (this is equivalent to requiring that the graph of f be type-definable). Assume furthermore that f is injective. Then* $\mathrm{Cb}' = f \circ \mathrm{Cb}$ *is another uniform canonical base map. Moreover, every uniform canonical base map can be obtained from any other in this manner.*

Proof. The main assertion follows from the fact that if f is definable and injective and $d\varphi(\bar{x}, Z)$ is a formula then $d\varphi(\bar{x}, f^{-1}(W))$ is also definable by a formula on the image of f. For the moreover part, given two uniform canonical base maps Cb and Cb$'$, the graph of the map $f\colon \mathrm{Cb}(p) \mapsto \mathrm{Cb}'(p)$ is type-definable (one canonical base has to give rise to the same definitions as the other, and this is a type-definable condition), so f is definable. $\blacksquare_{1.4}$

Thus, in the same way that a canonical base for a type is exactly anything which is inter-definable with another canonical base for that type, a uniform canonical base is exactly anything which is uniformly inter-definable with another uniform canonical base. A consequence of this (and of existence of uniformly canonical bases) is that in results such as the following the choice of uniform canonical bases is of no importance.

Lemma 1.5. *Let $\bar{z} = f(\bar{x}, \bar{y})$ be a definable function in T (say without parameters), possibly partial, and let* Cb *be uniform. Then the map* $f^{\mathrm{Cb}}\big(\mathrm{Cb}(\bar{a}/M), \bar{b}\big) = \mathrm{Cb}\big(f(\bar{a}, \bar{b})/M\big)$ *is definable as well for* $(\bar{a}, \bar{b}) \in \operatorname{dom} f$, $\bar{b} \in M$*, uniformly across all models of T. In case f is definable with parameters in some set A, so is f^{Cb}, uniformly across all models containing A.*

Proof. For the first assertion, it is enough to observe that we can define $\operatorname{tp}\big(f(\bar a)/M\big)$ by

$$\varphi\big(f(\bar a,\bar b),\bar c\big) \quad = \quad d\psi\big(\bar b\bar c,\operatorname{Cb}(\bar a/M)\big),$$

where $\psi(\bar x,\bar y\bar z)=\varphi\big(f(\bar x,\bar y),\bar z\big)$. The case with parameters is merely a special case. $\blacksquare_{1.5}$

Lemma 1.6. *Let* Cb *be a uniform canonical base map, say on the sort of n-tuples, into some infinite sort, and let $\operatorname{Cb}(p)_i$ denote its ith coordinate. Then the map $\bar a\mapsto\operatorname{Cb}(\bar a/M)_i$ is uniformly continuous, and uniformly so regardless of M.*

Proof. For a uniform canonical base map constructed from a uniform definition as discussed before Lemma 1.2 this follows from the fact that formulae are uniformly continuous. General case follows using Lemma 1.4 and the fact that definable functions are uniformly continuous. $\blacksquare_{1.6}$

Remark 1.7. The notion of a uniform canonical base map can be extended to simple theories, and the same results hold. Of course, canonical bases should then be taken in the sense of Hart, Kim and Pillay [HKP00], and one has to pay the usual price of working with hyper-imaginary sorts.

Now the question we asked earlier becomes

Question 1.8. Let T be a stable theory. Find a *natural* uniform canonical base map for T. In particular, one may want the image to be in the home sort, or in a restricted family of imaginary sorts.

Usually we shall aim for the image to lie in the home sort, plus the sort $\{T,F\}$ in the case of classical logic, or $[0,1]$ in the case of continuous logic.

Example 1.9. Let $T=IHS$, the theory of infinite dimensional Hilbert spaces, or rather, of unit balls thereof (from now on we shall tacitly identify Banach space structures with their unit balls).

The "folklore" canonical base for a type $p=\operatorname{tp}(\bar v/E)$ is the orthogonal projection $P_E(\bar v)$. This is not a uniform canonical base since it lacks enough information to recover p. On the other hand, it is obtained uniformly from any uniform canonical base of p, and by adding the missing information (in the sort $[-1,1]$) we obtain a uniform canonical base:

$$\operatorname{Cb}(\bar v/E)=\big(P_E(v_i),\langle v_i,v_j\rangle\big)_{i,j<n}.$$

This example, where we take a canonical base which is not uniform and make it uniform merely by adding information in a constant sort (namely, $\{T,F\}$ in classical logic, or $[0,1]$ in continuous logic) is a special case of the following.

Definition 1.10. Say that a canonical base map is *weakly uniform* if it can be obtained from a uniform map by composition with a definable function (which need not necessarily be injective, so the resulting canonical base need not suffice to recover the type uniformly – compare with Lemma 1.4).

For example, in the case of Hilbert spaces discussed above, the canonical base map $\operatorname{tp}(\bar v/E)\mapsto P_E(\bar v)$ is weakly uniform.

Proposition 1.11. *Let* Cb *be any uniform canonical base map, and let us write its target sort as $Z_0\times Z_1$, where Z_1 is a power of the constant sort. Let Cb_0 be the restriction to the sort Z_0. Then Cb_0 is a weakly uniform canonical base map. Conversely, every weakly uniform canonical base map can be obtained in this fashion.*

Proof. The main assertion is quite immediate, and it is the converse which we need to prove. Let Cb_0 be a weakly uniform canonical base map on a sort $\bar x$, with target sort Z_0. By definition, it is of the form $f\circ\operatorname{Cb}'$, where Cb' is a uniform canonical base map with target sort W and $f\colon \operatorname{img}\operatorname{Cb}'\to Z_0$ is definable. Let Φ be the set of all formulae $\varphi(\bar x,W)$. For every such formula, the value $\varphi\big(\bar a,\operatorname{Cb}'(\bar a/M)\big)$ is uniformly definable from $\operatorname{Cb}'(\bar a/M)$, call it $g_\varphi\big(\operatorname{Cb}'(\bar a/M)\big)$, and let $g=(g_\varphi)_{\varphi\in\Phi}$. Then $(f,g)\colon \operatorname{img}\operatorname{Cb}'\to Z_0\times Z_1$ is definable, and Z_1 is a power of the constant sort. If we show that (f, g) is injective then, by Lemma 1.4, we may conclude that $\mathrm{Cb} = (f, g) \circ \mathrm{Cb}'$ is the desired uniform canonical base map.

So let us consider a model $\mathcal{M}$ and two tuples $\bar{a}$ and $\bar{b}$ in the sort $\bar{x}$, lying in some extension $\mathcal{N} \succeq \mathcal{M}$. Let $C = \mathrm{Cb}'(\bar{a}/M)$, $D = \mathrm{Cb}'(\bar{b}/M)$, and assume that $(f, g)(C) = (f, g)(D)$. Then $g(C) = g(D)$ means that $\bar{a}C \equiv \bar{b}D$. Since also $f(C) = f(D)$, we have $\bar{a}f(C) \equiv \bar{b}f(C)$, i.e., $\bar{a} \equiv_{f(C)} \bar{b}$. Finally, by hypothesis, $f(C)$ is a canonical base for both types, whence $\bar{a} \equiv_M \bar{b}$ and therefore $C = D$. This completes the proof. $\blacksquare_{1.11}$

Thus our question can be restated as

Question 1.12. Let T be a stable theory. Find a natural weakly uniform canonical base map for T with image in the home sort.

Unfortunately, the canonical bases mentioned above for probability algebras and L_p lattices are not even weakly uniform, so we cannot apply Proposition 1.11 and the problem of finding uniform canonical bases requires some new ideas.

2. Uniform canonical bases in atomless probability spaces

The easier of the two "interesting cases" is that of atomless probability algebras. It does seem, however, that no uniform canonical bases exist in the home sort (of events), and that one must work instead in the (imaginary) sort of $[0, 1]$-valued random variables. This essentially boils down to working entirely within the theory *ARV* of atomless spaces of $[0, 1]$-valued random variables described in [Benb]. It is $\aleph_0$-stable, eliminates quantifiers, and admits definable continuous calculus: if $\tau \colon [0, 1]^n \to [0, 1]$ is any continuous function then the map $\bar{X} \mapsto \tau(\bar{X})$ is definable.

Fact 2.1. *Let $\bar{X} = X_0, \ldots, X_{n-1}$ be a tuple of bounded random variables. Then their joint distribution is determined by the sequence of moments $\big(\mathbf{E}[\bar{X}^{\bar{k}}]\big)_{\bar{k}\in\mathbf{N}^n}$ where $\bar{X}^{\bar{k}} = \prod X_i^{k_i}$.*

Similarly, their joint conditional distribution over a σ-algebra $\mathscr{B}$ is determined by the sequence $\big(\mathbf{E}[\bar{X}^{\bar{k}}|\mathscr{B}]\big)_{\bar{k}\in\mathbf{N}^n}$.

For a model $\mathcal{M} \vDash ARV$ let us write $\mathbf{E}[X|M]$ for $\mathbf{E}[X|\sigma(M)]$, which is itself a member of M.

Lemma 2.2. *Let $\mathcal{M} \preceq \mathcal{N} \vDash ARV$ and let $\bar{X} \in N^n$, $\bar{k} \in \mathbf{N}^n$. Then $\mathbf{E}[\bar{X}^{\bar{k}}|M]$ is uniformly definable from $\mathrm{Cb}(\bar{X}/M)$.*

Proof. By the definable continuous calculus, the function $(\bar{X}, Y) \mapsto |\bar{X}^{\bar{k}} - Y|^2$ is uniformly definable, and by Lemma 1.5 the predicate $\|\bar{X}^{\bar{k}} - y\|_2$ is uniformly definable for $y \in M$ from $\mathrm{Cb}(\bar{X}/M)$. For $Y \in M$ we have

$$Y = \mathbf{E}[\bar{X}^{\bar{k}}|M] \qquad \iff \qquad \|\bar{X}^{\bar{k}} - Y\|_2 = \inf_y \|\bar{X}^{\bar{k}} - y\|_2,$$

where the infimum is taken in $\mathcal{M}$. Thus the graph of the function $\mathrm{Cb}(\bar{X}/M) \mapsto \mathbf{E}[\bar{X}^{\bar{k}}|M]$ is type-definable in $\mathcal{M}$, whence it follows that the function itself is definable, and uniformly so in all models of *ARV*. $\blacksquare_{2.2}$

Theorem 2.3. *For n-types over models in ARV,*

$$\mathrm{Cb}(\bar{X}/M) = \big(\mathbf{E}[\bar{X}^k|M]\big)_{\bar{k}\in\mathbf{N}^n}$$

is a uniform canonical base (in the home sort).

Proof. By Lemma 2.2, this tuple is uniformly definable from any other uniform canonical base, so by Lemma 1.4 all that is left to show is that this tuple determines the type. By Fact 2.1, it determines the joint conditional distribution of $\bar{X}$ over $\sigma(M)$, which indeed determines $\mathrm{tp}(\bar{X}/M)$ by quantifier elimination. $\blacksquare_{2.3}$

3. Uniform canonical bases of in atomless L_p lattices

Recall that LpL denotes the theory of L_p lattices for some fixed $p \in [1, \infty)$, and that $ALpL$ denotes the theory of atomless ones. Stability, independence and related notions were studied for $ALpL$ by Berenstein, Henson and the author in [BBHb]. The theory $ALpL$ was shown to be $\aleph_0$-stable, and canonical bases of 1-types were described as tuples of *conditional slices* in the home sort (see Section 5 there). Even though they are very natural invariants of a 1-type, conditional slices are not uniform, or even weakly uniform, in the sense of the present paper. Our aim here is to replace the conditional slices with a related object which does provide a uniform canonical base.

We start by quickly recalling the *Krivine calculus* on Banach lattices (see also [LT79]).

Lemma 3.1. *Every lattice term $t(\bar{x})$ defines a function $t\colon \mathbf{R}^n \to \mathbf{R}$ which is finitely piecewise affine, continuous, and $\mathbf{R}^+$-homogeneous of degree one, by which we mean that $t(\alpha\bar{x}) = \alpha t(\bar{x})$ for all $\alpha \geq 0$.*

In addition, an arbitrary function $\varphi\colon \mathbf{R}^n \to \mathbf{R}$ is continuous and $\mathbf{R}^+$-homogeneous of degree one if and only if it can be approximated by lattice terms uniformly on every compact.

Proof. The first assertion, as well as the if part of the second, are clear. For the only if, let us assume that φ is continuous and $\mathbf{R}^+$-homogeneous of degree one. Then it is determined by its restriction to the unit sphere. Since every uniform approximation of φ on the unit sphere yields a uniform approximation on the entire unit ball, it will be enough to show that lattice terms are dense in $\mathcal{C}(S^{n-1}, \mathbf{R})$. They obviously form a lattice there, so it will be enough to show that for every distinct $\bar{x}, \bar{y} \in S^{n-1}$, every $a, b \in \mathbf{R}$ and every $\varepsilon > 0$, there is a lattice term t such that $|t(\bar{x}) - a|, |t(\bar{y}) - b| < \varepsilon$.

We may assume that $x_0 \neq y_0$. If $|x_0| = |y_0|$ then we may assume that $x_0 < 0 < y_0$ and define $t(\bar{z}) = \frac{a}{y_0} z_0^- + \frac{b}{y_0} z_0^+$. Otherwise, we may assume that $|x_0| < |y_0|$, in which case the opposite inequality must hold for some other coordinate, say $|y_1| < |x_1|$. Then $x_1 y_0 - x_0 y_1 \neq 0$ and we may define

$$t(\bar{z}) = \frac{b x_1 - a y_1}{x_1 y_0 - x_0 y_1} z_0 + \frac{a y_0 - b x_0}{x_1 y_0 - x_0 y_1} z_1.$$

Either way, $t(\bar{x}) = a$ and $t(\bar{y}) = b$, which is even better than what we needed. $\blacksquare_{3.1}$

It is also a fact that if t is a lattice term with bound $\|t\|$ on $[-1, 1]^n$ then for any Banach lattice E and every sequence $\bar{f} \in E^n$ one has $|t(\bar{f})| \leq \|t\| \bigvee_i |f_i|$. It follows that if $t_k \to \varphi$ uniformly on $[-1, 1]^n$ then $t_k(\bar{f})$ converges in norm to a limit $\varphi(\bar{f})$ which does not depend on the choice of converging lattice terms, and at a rate which only depends on the sequence t_k and on $\| \bigvee_i |f_i| \|$. It follows that the map $\varphi\colon E^n \to E$ is uniformly definable across all Banach lattices.

Convention 3.2. For $\alpha > 0$ we extend $x \mapsto x^\alpha$ to the whole real line by $(-x)^\alpha = -x^\alpha$ (so $(-7)^2 = -49$).

Lemma 3.3. *For every $p, q \in [1, \infty)$ the theories LpL and LqL are quantifier-free bidefinable. More exactly, if $E = L_p(\Omega)$ and $F = L_q(\Omega)$ then we may identify their respective underlying sets via the bijection $f \mapsto f^{\frac{p}{q}}$. Under this identification, each Banach lattice structure is quantifier-free definable in the other.*

Moreover, if $q \in (1, \infty)$ and $q' = \frac{q}{q-1}$ is its conjugate exponent, then the duality pairing $\langle f, g \rangle = \int fg$ on $L_q \times L_{q'}$ is definable in L_p.

Proof. Immediate. $\blacksquare_{3.3}$

For the time being we consider a 1-type over a model $\operatorname{tp}(f/E)$.

Lemma 3.4. *Assume that $p > 1$, and let* Cb *be any uniform canonical base map. Then $\operatorname{Cb}(f/E) \mapsto \mathbf{E}[f|E]$ is definable in $ALpL$. More generally, if $t(x, \bar{y})$ is any lattice term, then the map $\big(\operatorname{Cb}(f/E), \bar{g}\big) \mapsto \mathbf{E}\big[t(f, \bar{g})|E\big]$, where $\bar{g} \in E$, is definable.*

Proof. Since $\|\mathbf{E}[f|E]\| \leq \|f\|$ for all f, it is enough to show that the graph of the map is uniformly type-definable in the canonical base. Indeed, the function $(x, y) \mapsto x^{\frac{1}{p}} y^{\frac{p-1}{p}}$ is continuous and $\mathbf{R}^+$-homogeneous of degree one. Therefore $(h, g) \mapsto \|h^{\frac{1}{p}} g^{\frac{p-1}{p}}\|^p = \int |h||g|^{p-1}$ is a definable predicate. Separating into positive and negative parts we see that $\varphi(x, z) = \int xz^{p-1}$ is a definable predicate as well. Then in E we have $x = \mathbf{E}[f|E]$ if and only if

$$\sup_z \left|\varphi(x, z) - d\varphi\big(z, \mathrm{Cb}(f/E)\big)\right| = 0,$$

and this remains true if we restrict z to the unit ball. This concludes the proof. The case of a lattice term follows by Lemma 1.5. $\blacksquare_{3.4}$

Remark 3.5. It follows that at least for $p > 1$, the conditional expectation with respect to a sub-lattice is intrinsic to the structures and does not depend on any presentation as concrete L_p spaces. Using similar techniques (namely, interpreting an L_q lattice in an L_p lattice) we shall see later on that the same follows for $p = 1$, although without the uniformity proved above. This fact has already been shown in [BBHb] for all p but without uniformity.

Let us consider the map $g \mapsto \mathbf{E}[g \dot{-} f|E]$ on E. It is clearly determined by $\mathrm{tp}(f/E)$, and it follows from Lemma 3.4 that for $p > 1$ it is even uniformly definable from the canonical base. It follows from what we do later that this map also contains sufficient information in order to recover the type $\mathrm{tp}(f/E)$. We can code the map (and thus the type, and uniformly so when $p > 1$) by a canonical parameter for the predicate $d\big(x, \mathbf{E}[y \dot{-} f|E]\big)$, but as we explained in the introduction this is exactly what we wish to avoid doing.

Instead, we observe that the map $g \mapsto \mathbf{E}[g \dot{-} f|E]$ is convex on E. It is therefore equal to its double Legendre transform (with respect to the multiplication $E \times \mathbf{R} \to E$), and in particular can be recovered from its Legendre transform, which we shall denote by $t \mapsto \mathbf{E}_t[f|E]$:

$$\mathbf{E}_t[f|E] = \sup_{g \in E} tg - \mathbf{E}[g \dot{-} f|E].$$

We observe that $\mathbf{E}_t[f|E] = +\infty$ for $t \notin [0, 1]$, and we claim that $\mathbf{E}_t[f|E] \in E$ for $t \in [0, 1]$. Indeed, it is also not difficult to see that $\mathbf{E}_0[f|E] = 0$ and $\mathbf{E}_1[f|E] = \mathbf{E}[f|E]$. Since $t \mapsto \mathbf{E}_t[f|E]$ is convex (as a Legendre transform) we obtain $\mathbf{E}_t[f|E] \leq t\mathbf{E}[f|E]$ for $0 < t < 1$, as desired. The double Legendre transform is then

$$\mathbf{E}[g \dot{-} f|E] = \sup_{t \in [0,1]} tg - \mathbf{E}_t[f|E].$$

In order to see a little clearer, let us fix a concrete presentation $E = L_p(X)$, and let $E' = L_p\Big(\big(X \times [0, 1]\big) \cup [0, 1]\Big)$, where the intervals are taken with the Lebesgue measure. Then E embeds canonically in E', and by quantifier elimination for $ALpL$ the embedding is elementary. It is shown in [BBHb] that every 1-type over E admits an "increasing" realisation in E'. We may therefore assume that $f \in E'$, and that for each $x \in X$, the map $t \mapsto f(x, t)$ is increasing. In this case, $\mathbf{E}[f|E] \in E$ is merely the map $x \mapsto \int_0^1 f(x, t)\, dt$. For $t \in (0, 1)$ and $x \in X$ let $f_t(x) = f(x, t)$. Then the map $t \mapsto f_t$ is increasing, and we have (where $\|\cdot\| = \|\cdot\|_p$)

$$\|f\|^p \geq \|f{\restriction}_{X \times [0,1]}\|^p = \int_0^1 \|f_t\|^p\, dt = \int_0^1 \|f_t^+\|^p + \|f_t^-\|^p\, dt.$$

Since $\|f_t^+\|$ (respectively, $\|f_t^-\|$) is increasing (respectively, decreasing) in t we get

$$\|f_t^+\|^p \leq \frac{\|f\|^p}{1 - t}, \qquad \|f_t^-\|^p \leq \frac{\|f\|^p}{t},$$

whereby

$$\|f_t\| \leq \frac{\|f\|}{(t-t^2)^{1/p}}.$$

In particular, $f_t \in E$ for all $t \in (0,1)$ (in fact one can actually get $\|f_t\| \leq \|f\| \min(t, 1-t)^{-1/p}$). It is not difficult to see now that

$$tf_t - \mathbf{E}[f_t \dot{-} f|E] = \mathbf{E}_t[f|E] = \int_0^t f_s\, ds.$$

The first equality just means that the supremum in the definition of $\mathbf{E}_t[f|E]$ is attained at f_t.

In fact, it is a general fact that if F is convex and $F^*(t) = \sup_x tx - F(x)$ is its Legendre transform, then the supremum is attained at x if and only $D_t^- F^*(t) \leq x \leq D_t^+ F^*(t)$, where $D^\pm$ denote the derivatives on the left and on the right, respectively. In our case we have $F(g) = \mathbf{E}[g \dot{-} f|E]$ and $F^*(t) = \mathbf{E}_t[f|E]$, so the first equality above is equivalent to $D_t^- \mathbf{E}_t[f|E] \leq f_t \leq D_t^+ \mathbf{E}_t[f|E]$ for all t, which in turn is equivalent to the second. A comparison with [BBHb] yields that f_t is equal to the conditional slice $\mathbf{S}_{1-t}(f/E)$ for all t where the two one-sided derivatives agree, and in particular for almost all t. Notice that this description of f_t in terms of one-sided derivatives also means that it f_t is a well-determined member of E for almost all t in a manner which is intrinsic to the type of f over E.

Lemma 3.6. *For each t, Assume that $p > 1$. Then $\mathrm{Cb}(f/E) \mapsto \mathbf{E}_t[f|E]$ is definable in ALpL for all $t \in [0,1]$.*

Proof. For $t = 0, 1$ this is already known, so we may assume that $t \in (0,1)$. It follows from the definition that in E:

$$\big\|x - \mathbf{E}_t[f|E]\big\|^p = \inf_y \Big\|\big(x - ty + \mathbf{E}[y \dot{-} f|E]\big)^+\Big\|^p + \sup_y \Big\|\big(x - ty + \mathbf{E}[y \dot{-} f|E]\big)^-\Big\|^p.$$

It is enough to restrict the quantifiers on y to $\|y\| \leq \|f_t\| \leq \frac{\|f\|}{\sqrt[p]{t-t^2}}$. Together with Lemma 3.4, this means that $d\big(x, \mathbf{E}_t[f|E]\big)$ is uniformly definable from $\mathrm{Cb}(f/E)$, and the proof is complete. $\blacksquare_{3.6}$

Theorem 3.7. *For every $p \in (1,\infty)$ and every dense subset $D \subseteq (0,1)$ (e.g., $D = \mathbf{Q} \cap (0,1)$), the tuple $\big(\|f^+\|, \|f^-\|, \mathbf{E}_t[f|E]\big)_{t \in D}$ is a uniform canonical base for $\mathrm{tp}(f/E)$, and $\big(\mathbf{E}_t[f|E]\big)_{t\in D}$, is a weakly uniform canonical base in the home sort.*

Proof. By Proposition 1.11 it is enough to prove the first assertion. We have already seen that $\mathrm{Cb}(f/E) \mapsto \mathbf{E}_t[f|E]$ is definable, and clearly $\mathrm{Cb}(f/E) \mapsto \|f^\pm\|$ are, so by Lemma 1.4 all that is left is to show that the tuple $\big(\|f^+\|, \|f^-\|, \mathbf{E}_t[f|E]\big)_{t\in D}$ determines $\mathrm{tp}(f/E)$.

Since $t \mapsto \mathbf{E}_t[f|E]$ is convex, it is determined on a dense subset, so for all $t \in (0,1)$ we may define $g_t = D_t^- \mathbf{E}_t[f|E] \in E$. Working again in E' defined as above, define $g(x,t) = g_t(x)$. Then $g\restriction_{X\times[0,1]} = f_{X\times[0,1]}$ (almost everywhere, and therefore in E'). In particular, $\|g^\pm\restriction_{X\times[0,1]}\| \leq \|f^\pm\|$, and we may define g on the disjoint copy of $[0,1]$ so that $\|g^\pm\| = \|f^\pm\|$. Then $g \vDash \mathrm{tp}(f/E)$, and the proof is complete. $\blacksquare_{3.7}$

Notice that $\mathbf{E}_t[f|E] \to 0$ as $t \to 0$, by dominated convergence, and similarly $\mathbf{E}_t[f|E] \to \mathbf{E}[f|E]$ as $t \to 1$. Moreover, for $p > 1$ (fixed) the rate of convergence depends uniformly on $\|f\|$. Indeed, otherwise Lemma 3.6 together with a compactness argument would yield a type (or a canonical base of a type, which is the same thing) for which convergence fails altogether. On the other hand, for $p = 1$, consider for some ε the case where $\mu(X) = 1$ and $f_\varepsilon(x,t) = -\varepsilon^{-1}\mathbf{1}_{X\times[0,\varepsilon]}$. Then $\|f_\varepsilon\| = 1$ and $\mathbf{E}_\varepsilon[f_\varepsilon|E] = -\mathbf{1}_X$, also of norm one, so the rate of convergence is not uniform. Thus Lemma 3.6, and therefore Lemma 3.4, fail for $p = 1$. This is essentially the only obstacle, and by keeping away from the endpoints of $[0,1]$ we do manage to get an analogue of Theorem 3.7 for $p = 1$. For $0 \leq t < s \leq 1$ let us define

$$\mathbf{E}_{[t,s]}[f|E] = \mathbf{E}_s[f|E] - \mathbf{E}_t[f|E] = \int_t^s f_r \, dr.$$

Lemma 3.8. *Let $p \in [1,\infty)$. Then for every $0 < s < t < 1$, the map $\mathrm{Cb}(f/E) \mapsto \mathbf{E}_{[t,s]}[f|E]$ is definable.*

Proof. For $p > 1$ this is already known, so we only need to deal with the case of $p = 1$. For each $q > 1$ we may apply the bidefinability of $AL1L$ and $ALqL$, and calculate $\mathbf{E}_{[t,s]}[f|E]^{L_q} = \mathbf{E}_{[t,s]}[f^{1/q}|E]^q$ uniformly from $\mathrm{Cb}(f/E)$. It will be enough to show that as $q \to 1$, $\mathbf{E}_{[t,s]}[f|E]^{L_q} \to \mathbf{E}_{[t,s]}[f|E]$ at a rate which only depends on $\|f\|$. We may choose a concrete representation where $h = \mathbf{E}\big[|f|\big|E\big]$ is an indicator function, and consider what happens in a single fibre over E. There all functions in E are constants, and we may identify f with the function $f(t) = f_t$, which is increasing on $[0,1]$. If $h = 0$ then everything is zero, so we may assume that $h = \int_0^1 |f| = 1$. Since f is increasing, we must have $f(r) \in [-\frac{1}{t}, \frac{1}{1-s}]$ for all $r \in [t,s]$, and $\int_t^s f(r)^{1/q} \, dr \in [-\frac{1}{t}, \frac{1}{1-s}]$ as well. It follows that for any desired $\varepsilon > 0$ there exists $q_0 > 1$, depending only on t, s and ε, such that for all $1 < q < q_0$:

$$\left| \left(\int_t^s f(r)^{\frac{1}{q}} \, dr \right)^q - \int_t^s f(r) \, dr \right| < \varepsilon.$$

Integrating over all fibres we obtain

$$\left| \mathbf{E}_{[t,s]}[f|E]^{L_q} - \mathbf{E}_{[t,s]}[f|E] \right| < \varepsilon \|f\|,$$

as desired. $\blacksquare_{3.8}$

Theorem 3.9. *For every $p \in [1,\infty)$ and every dense subset $D \subseteq (0,1)$ (e.g., $D = \mathbf{Q} \cap (0,1)$), the tuple $\big(\|f^+\|, \|f^-\|, \mathbf{E}_{[t,s]}[f|E]\big)_{t,s \in D, t<s}$ is a uniform canonical base for $\mathrm{tp}(f/E)$, and $\big(\mathbf{E}_{[t,s]}[f|E]\big)_{t,s \in D, t<s}$ is a weakly uniform canonical base in the home sort.*

Proof. Same argument as for Theorem 3.7. $\blacksquare_{3.9}$

We have thus produced (weakly) uniform canonical bases in the home sort for 1-types in $ALpL$. For n-types, we use the following general fact. Recall first that if $E \subseteq E'$ are two L_p lattices then each member of E' can be written uniquely as $f = f{\restriction}_E + f{\restriction}_{E^\perp}$, where $f{\restriction}_{E^\perp}$ is orthogonal to E and $f{\restriction}_E$ is orthogonal to $E^\perp = \{g \in E' : g \perp E\}$. If $E = L_p(X,\Sigma,\mu) \subseteq L_p(X',\Sigma',\mu')$ (where $(X,\Sigma,\mu) \subseteq (X',\Sigma',\mu')$ is an extension of measure spaces) then $f{\restriction}_E = f\mathbf{1}_X$ and $f{\restriction}_{E^\perp} = f\mathbf{1}_{X' \smallsetminus X}$. In particular, $\cdot{\restriction}_E$ and $\cdot{\restriction}_{E^\perp}$ are linear lattice homomorphisms.

Fact 3.10. *The n-type $\mathrm{tp}(\bar{f}{\restriction}_E/E)$ is determined by the 1-types $\mathrm{tp}(\bar{k} \cdot \bar{f}/E)$, where $\bar{k} \cdot \bar{f} = \sum_i k_i f_i$ and $\bar{k}$ varies over $\mathbf{Z}^n$.*

Proof. Indeed, this information determines $\mathrm{tp}(\bar{t} \cdot \bar{f}/E)$, and in particular $\mathrm{tp}(\bar{t} \cdot \bar{f}{\restriction}_E/E)$ for all $\bar{t} \in \mathbf{Q}^n$ and therefore for all $\bar{t} \in \mathbf{R}^n$. Now apply [BBHb, Proposition 3.7]. $\blacksquare_{3.10}$

Theorem 3.11. *Let Cb be a uniform canonical base map for 1-types. Then*

$$\mathrm{tp}(\bar{f}/E) \mapsto \big(\mathrm{Cb}(\bar{k} \cdot \bar{f}/E), \mathrm{tp}(\bar{f})\big)_{\bar{k} \in \mathbf{Z}^n}$$

is a uniform canonical base map for n-types.

If Cb is a weakly uniform canonical base map for 1-types then

$$\mathrm{tp}(\bar{f}/E) \mapsto \big(\mathrm{Cb}(\bar{k} \cdot \bar{f}/E)\big)_{\bar{k} \in \mathbf{Z}^n}$$

is a weakly uniform canonical base map for n-types.

We may view $\mathrm{tp}(\bar{f})$ as a sequence in $[0,1]$ via any embedding of $\mathrm{S}_n(ALpL)$ in $[0,1]^{\aleph_0}$.

Proof. For the first assertion, it is enough to show that $\mathrm{tp}(\bar f/E)$ is determined by this data. Indeed, $\mathrm{tp}(\bar f\restriction_E/E)$ is already known to be determined. Let $h = \sum_i |f_i|$. Then $\mathrm{tp}(\bar f\restriction_E/E)$ determines $\mathrm{tp}(h\restriction_E/E)$, and in particular $\|h\restriction_E\|$, while $\|h\|$ is determined by $\mathrm{tp}(\bar f)$. Thus $\|h\restriction_{E^\perp}\|$ is known. Alongside the facts that $h\restriction_{E^\perp}$ is positive and orthogonal to E, this is enough to determine $\mathrm{tp}(h\restriction_{E^\perp}/E)$. Since in any case we are only interested in the type over E, we may assume that $h\restriction_E$ and $h\restriction_{E^\perp}$ are known. We may further assume that $h\restriction_E = \mathbf{1}_A$ and $h\restriction_{E^\perp} = \mathbf{1}_B$ in some concrete presentation of the ambient space. Now $\mathrm{tp}(\bar f)$ determines $\mathrm{tp}(\bar f/h)$, which, again by [BBHb, Proposition 3.7], can be identified with the joint conditional distribution of $\bar f$ with respect to $\{\varnothing, A\cup B\}$ (which is essentially the same thing as the distribution of $\bar f$ restricted to $A\cup B$, with the caveat that $A\cup B$ has finite measure which is not necessarily one, i.e., is not necessarily a probability space). Similarly, $\mathrm{tp}(\bar f\restriction_E/E)$ determines $\mathrm{tp}(\bar f\restriction_E)$ and thus $\mathrm{tp}(\bar f\restriction_E/h\restriction_E)$, which can be identified with the joint conditional distribution of $\bar f\restriction_E$ with respect to $\{\varnothing, A\}$. Subtracting, we obtain the joint conditional distribution of $\bar f\restriction_{E^\perp}$ with respect to $\{\varnothing, B\}$, namely $\mathrm{tp}(\bar f\restriction_{E^\perp}/h\restriction_{E^\perp})$, and thus $\mathrm{tp}(\bar f\restriction_{E^\perp})$ and finally $\mathrm{tp}(\bar f\restriction_{E^\perp}/E, \bar f\restriction_E)$. Thus $\mathrm{tp}(\bar f/E)$ is known and the proof of the first assertion is complete.

For the second assertion it is enough to show that the tuple is a canonical base for the type, i.e., is inter-definable with some other canonical base for the type, a fact which follows from the first assertion. $\blacksquare_{3.11}$

Corollary 3.12. *For every $p \in [1,\infty)$ and every dense subset $D \subseteq (0,1)$ the tuple $\big(\mathrm{tp}(\bar f), \mathbf{E}_{[t,s]}[\bar k \cdot \bar f|E]\big)_{t,s\in D, t<s, \bar k\in \mathbf{Z}^n}$ is a uniform canonical base for $\mathrm{tp}(\bar f/E)$, and $\big(\mathbf{E}_{[t,s]}[\bar k\cdot \bar f|E]\big)_{t,s\in D,t<s,\bar k\in\mathbf{Z}^n}$ is a weakly uniform canonical base in the home sort. When $p > 1$ we may replace $\mathbf{E}_{[t,s]}$ with $\mathbf{E}_t$.*

Remark 3.13. At least for $p = 1$ this cannot be improved, in the sense that the types of every $\bar k \cdot \bar f$ need not determine $\mathrm{tp}(\bar f)$. Indeed, let f_1, f_2 and f_3 be disjoint positive functions of norm one, and let

$$g = f_1 - f_2, \qquad h = f_1 + f_2 - 2f_3.$$

Then $\mathrm{tp}(kg + \ell h) = \mathrm{tp}(kg - \ell h)$ for all k, ℓ, but $\mathrm{tp}(g, h) \neq \mathrm{tp}(g, -h)$.

4. On uniform canonical bases and beautiful pairs

It is implicitly shown by Poizat [Poi83], based on Shelah's f.c.p. Theorem [She90], that a stable classical theory does not have the finite cover property if and only if the set of uniform canonical bases (for all types in any one given sort) is definable, rather than merely type-definable (here, a definable set in an infinite sort means a set which is closed under coordinate-wise convergence, and such that the projection to each finite sub-sort is definable in the ordinary sense). A similar result should hold for continuous logic, where the finite cover property (and in particular Shelah's f.c.p. Theorem) have not yet been properly studied. Here we concentrate on the relation between the existence of a good first order theory for beautiful pairs and the definability of the sets of uniform canonical bases.

We fix a stable theory T in a language $\mathcal{L}$ admitting quantifier elimination as well as a uniform canonical base map Cb. We may write the latter as $(\mathrm{Cb}_n)_n$, since it consists of a map for the sort of n-tuples for each n (we shall assume that $\mathcal{L}$ is single sorted, otherwise even more complex notation is required). We define $\mathcal{L}_P = \mathcal{L} \cup \{P\}$, where P is a new unary predicate symbol (1-Lipschitz, in the continuous setting). We also define $\mathcal{L}_{\mathrm{Cb}}$ to consist of $\mathcal{L}$ along with, for each n, n-ary function symbols to the target sorts of Cb_n. We denote the (possibly infinite) tuple of these new function symbols $f_{\mathrm{Cb}}(\bar x)$, where $n = |\bar x|$. In the continuous setting, uniform continuity moduli for the f_{Cb} are as per Lemma 1.6. We let $\mathcal{L}_{P,\mathrm{Cb}} = \mathcal{L}_P \cup \mathcal{L}_{\mathrm{Cb}}$.

By a *pair* of models of T we mean any elementary extension $\mathcal{N} \preceq \mathcal{M} \vDash T$. We shall identify such a pair with the structures $(\mathcal{M}, P)$, $(\mathcal{M}, f_{\mathrm{Cb}})$ or $(\mathcal{M}, P, f_{\mathrm{Cb}})$, as will be convenient, where $P(x) = d(x, N)$ and $f_{\mathrm{Cb}}(\bar x) = \mathrm{Cb}(\bar x/N)$. The property that the predicate P defines an elementary sub-structure is elementary, so the class of all pairs of models of T is elementary as well, of theory $T_{P,0}$. Similarly, $T_{P,\mathrm{Cb},0}$ will be the $\mathcal{L}_{P,\mathrm{Cb}}$-theory of pairs, which consists in addition of the axioms saying that $f_{\mathrm{Cb}}(\bar x) = \mathrm{Cb}(\bar x/P)$.

It is easy to check that all these axioms are indeed expressible by an inductive $\mathcal{L}_P$-theory and an inductive $\mathcal{L}_{P,\mathrm{Cb}}$-theory, respectively. Clearly, $T_{P,\mathrm{Cb},0}$ is a definitional expansion of $T_{P,0}$, so we may unambiguously refer to a model of $T_{P,\mathrm{Cb},0}$ as $(\mathcal{M}, P)$. On the other hand, the predicate P is also superfluous in $T_{P,\mathrm{Cb},0}$, since it can be recovered in the classical and continuous cases, respectively, as

$$P(x) = \exists y\, d[x = y]\big(y, \mathrm{Cb}(x/P)\big), \qquad P(x) = \inf_y\, d[d(x,y)]\big(y, \mathrm{Cb}(x/P)\big).$$

Since T admits quantifier elimination, the formulae on the right hand side can be taken to be quantifier-free. We may therefore express the same properties as above in an inductive $\mathcal{L}_{\mathrm{Cb}}$-theory $T_{\mathrm{Cb},0}$, for which $T_{P,\mathrm{Cb},0}$ is merely a quantifier-free definitional expansion. We may therefore work quite interchangeably in one setting or the other, i.e., with or without the predicate P.

Lemma 4.1. *The theory $T_{\mathrm{Cb},0}$ admits amalgamation over arbitrary sets. If T is complete, then $T_{\mathrm{Cb},0}$ also admits the joint embedding property.*

Proof. Let $(\mathcal{M}_i, P) \vDash T_{\mathrm{Cb},0}$ for $i = 0, 1$ and let A be a common sub-structure. Then $\mathcal{M}_0 \equiv \mathcal{M}_1$ (even if T is incomplete) so let $\mathcal{N}$ be a large model of the common theory in which we seek to embed this configuration, at a first time as $\mathcal{L}$-structures. First, we may place A and $P(M_i)$ in $\mathcal{N}$ so that $P(M_0) \mathop{\smile\!\!\!\!\!\!\;|}_A P(M_1)$. Since A is a sub-structure in $\mathcal{L}_{\mathrm{Cb}}$, we know that $A \mathop{\smile\!\!\!\!\!\!\;|}_{P(A)} P(M_i)$. It follows that $A \mathop{\smile\!\!\!\!\!\!\;|}_{P(A)} P(M_0)P(M_1)$, and we may choose a sub-model $P(\mathcal{N}) \preceq \mathcal{N}$ such that $P(N) \supseteq P(M_i)$ and $A \mathop{\smile\!\!\!\!\!\!\;|}_{P(A)} P(N)$. We now embed each M_i such that $M_i \mathop{\smile\!\!\!\!\!\!\;|}_{A,P(M_i)} P(N)$. Then in particular $M_i \mathop{\smile\!\!\!\!\!\!\;|}_{P(M_i)} P(N)$. It follows that the embeddings of $(\mathcal{M}_i, P)$ in $(\mathcal{N}, P)$ respect $\mathcal{L}_{\mathrm{Cb}}$, and we are done. If T is complete then we can amalgamate any two models of $T_{\mathrm{Cb},0}$ over the empty substructure. $\blacksquare_{4.1}$

(This argument already appears, in essence, in [BPV03], the only novelty is that we use the language $\mathcal{L}_{\mathrm{Cb}}$ to ensure that every sub-structure is P-independent, i.e., verifies $A \mathop{\smile\!\!\!\!\!\!\;|}_{P(A)} P$.)

It follows that a model companion of $T_{\mathrm{Cb},0}$ (or of $T_{P,\mathrm{Cb},0}$, this is the same thing), if it exists, eliminates quantifiers, i.e., it is a model completion. Even if it does not exist we may still consider it as a Robinson theory in the sense of Hrushovski [Hru97].

Lemma 4.2. *Modulo $T_{\mathrm{Cb},0}$, the restriction of every quantifier-free $\mathcal{L}_{\mathrm{Cb}}$-formula to P is $\mathcal{L}$-definable there.*

Proof. This follows immediately from the fact that the map $\bar{a} \mapsto \mathrm{Cb}(\bar{a}/\bar{a})$ is definable in $\mathcal{L}$. $\blacksquare_{4.2}$

Definition 4.3. Following Poizat [Poi83], we say that a pair $(\mathcal{M}, P)$ of models of T is *beautiful* if P is approximately $\aleph_0$-saturated and $\mathcal{N}$ is approximately $\aleph_0$-saturated over $\mathcal{M}$. We define T^b_{Cb} to be the $\mathcal{L}_{\mathrm{Cb}}$-theory of all beautiful pairs of models of T.

Theorem 4.4. *Let T be a stable theory with quantifier elimination and a uniform canonical base map, as above. Then the following are equivalent.*

(i) *The image set $\operatorname{img}\mathrm{Cb}_n$ is definable for each n (i.e., its projection to every finite sort is a definable set in that sort).*

(ii) *If $(\mathcal{M}, P) \vDash T^b_{\mathrm{Cb}}$ is κ-saturated for some $\kappa > |T|$ then $\mathcal{M}$ is κ-saturated over $P(M)$. In particular, every sufficiently saturated model of T^b_{Cb} is itself a beautiful pair.*

(iii) *The theory $T_{\mathrm{Cb},0}$ admits a companion T_{Cb} such that for some $\kappa > |T|$, if $(\mathcal{M}, P) \vDash T_{\mathrm{Cb}}$ is κ-saturated then $\mathcal{M}$ is κ-saturated over $P(M)$. Moreover, such a companion is necessarily the model companion of $T_{\mathrm{Cb},0}$.*

(iv) *The theory $T_{\mathrm{Cb},0}$ admits a model completion T_{Cb} (i.e., its model companion exists and eliminates quantifiers).*

If T is complete, this is further equivalent to:

(v) *Let* $(\mathcal{M}, P) \vDash T_{\mathrm{Cb},0}$*, where* $P(\mathcal{M})$ *is* $|T|^+$*-saturated and* $\mathcal{M}$ *is* $|T|^+$*-saturated over* $P(M)$*. Then* $(\mathcal{M}, P)$ *is* $\aleph_0$*-saturated (and is, moreover, a model of the model companion).*

(vi) *There exists an approximately* $\aleph_0$*-saturated beautiful pair* $(\mathcal{M}, P)$*.*

Proof. (i) $\Longrightarrow$ (ii). Let T_{Cb} consist of $T_{\mathrm{Cb},0}$ along with the axioms saying that for every canonical base $Z \in \operatorname{img}\mathrm{Cb}$ and every tuple $\bar{w}$, the type defined by Z on the domain $P \cup \bar{w}$ is finitely realised (or, in the continuous setting, approximately finitely realised). Since the set $\operatorname{img}\mathrm{Cb}$ is definable, this axiom can indeed be expressed, and clearly $T_{\mathrm{Cb}} \subseteq T^b_{\mathrm{Cb}}$.

Now let $(\mathcal{M}, P) \vDash T_{\mathrm{Cb}}$ be κ-saturated, in which case $P(\mathcal{M})$ is κ-saturated as well. Let also $\bar{a}$ be a tuple in some elementary extension of $\mathcal{M}$, $C = \mathrm{Cb}(\bar{a}/M) \in M$, and let $A \subseteq M$, $|A| < \kappa$. Let $\pi(\bar{x})$ be the partial $\mathcal{L}_{\mathrm{Cb}}$-type saying that $\bar{x}$ realises $\mathrm{tp}(\bar{a}/A \cup P)$, i.e., that

$$\sup_{\bar{y}\in P}\big|\varphi(\bar{x}, \bar{y}\bar{b}) - d\varphi(\bar{y}\bar{b}, C)\big| = 0$$

for each formula $\varphi(\bar{x}, \bar{y}\bar{z})$ and $\bar{b} \in A$. By T_{Cb}, this partial type is approximately finitely realised in $(\mathcal{M}, P)$, and since $|A \cup C| < \kappa$, it is realised there.

(ii) $\Longrightarrow$ (iii). It is easy to check that $T_{\mathrm{Cb},0}$ and T^b_{Cb} are companions, whence the main assertion. For the moreover part, let T_{Cb} be any companion with the stated property, and we shall show that all its models are existentially closed. Indeed, let us consider an extension $(\mathcal{M}, P) \subseteq (\mathcal{N}, P)$ of models of T_{Cb}, and we need to show that $(\mathcal{M}, P)$ is existentially closed in $(\mathcal{N}, P)$. Since the latter is an elementary property of an extension, we may replace the extension with any elementary extension thereof (technically speaking, we represent the extension by $(\mathcal{N}, P, Q)$, where $Q(x) = d(x, M)$, and take an elementary extension of this). We may therefore assume that $(\mathcal{M}, P)$ is κ-saturated, so in particular $P(\mathcal{M})$ is κ-saturated, and by assumption $\mathcal{M}$ is κ-saturated over $P(M)$.

Now let $\bar{b} \subseteq M^m$, $\bar{a} \in N^n$, and let $p(\bar{x}, \bar{y}) = \mathrm{tp}\big(\bar{a}, \bar{b}/P(N)\big)$, $q(\bar{y}) = \mathrm{tp}\big(\bar{b}/P(N)\big)$, $C = \mathrm{Cb}(p)$, $D = \mathrm{Cb}(q)$. By saturation of $P(\mathcal{M})$ we may find $C' \subseteq P(M)$ such that $C' \equiv_D C$, and define $p'(\bar{x}, \bar{y})$ to be the type over $P(N)$ defined by C'. Now, $C \equiv_D C'$ along with $q \subseteq p$ yields $q \subseteq p'$, so $p'(\bar{x}, \bar{b})$ is consistent. Therefore, its restriction to $P(M), \bar{b}$ is realised in M, say by $\bar{a}'$. Then $\bar{a}$ and $\bar{a}'$ have the same quantifier-free $\mathcal{L}_{\mathrm{Cb}}$-type over $\bar{b}$, which concludes the proof.

(iii) $\Longrightarrow$ (iv). By Lemma 4.1, the model companion eliminates quantifiers.

(iv) $\Longrightarrow$ (i). The set $(\operatorname{img}\mathrm{Cb})^P$ is definable in T_{Cb}, as the image of a definable set (the entire universe) under a definable function. By quantifier elimination, it is quantifier-free definable. It follows that the set $(\operatorname{img}\mathrm{Cb})^P$ is $\mathcal{L}$-definable in P, which means exactly that $\operatorname{img}\mathrm{Cb}$ is definable.

(iv) $\Longrightarrow$ (v). A close inspection of the argument for (ii) $\Longrightarrow$ (iii) reveals that it proves the following intermediary result: if $(\mathcal{M}, P) \vDash T_{\mathrm{Cb},0}$, $P(\mathcal{M})$ is κ-saturated, and $\mathcal{M}$ is κ-saturated over $P(M)$, then $(\mathcal{M}, P)$ is $\aleph_0$-saturated. In addition, such a pair is clearly a model of T_{Cb} as given there.

(v) $\Longrightarrow$ (vi). Clear.

(vi) $\Longrightarrow$ (i). Here we also assume that T is complete. Assume that $\operatorname{img}\mathrm{Cb}_n$ is not definable, keeping in mind that it is always type-definable. This means that for some $\varepsilon > 0$, the set of types q in the sort of $\operatorname{img}\mathrm{Cb}_n$ satisfying $d(q, \operatorname{img}\mathrm{Cb}_n) \geq \varepsilon$ has an accumulation point in $\operatorname{img}\mathrm{Cb}_n$. By our saturation assumption for $P(\mathcal{M})$, the partial type saying that $Z \in (\operatorname{img}\mathrm{Cb}_n)^P$ and $d\big(Z, f_{\mathrm{Cb}}(M^n)\big) \geq \varepsilon$ is approximately finitely realised in $(\mathcal{M}, P)$, and therefore, by the saturation assumption for $(\mathcal{M}, P)$, realised there, say by C. Then the type $p \in \mathrm{S}_n(P)$ whose canonical base is C is (approximately) realised in $\mathcal{M}$, by the saturation assumption for $\mathcal{M}$, so $d\big(C, f_{\mathrm{Cb}}(M^n)\big) = 0$, a contradiction. $\blacksquare_{4.4}$

Corollary 4.5. *The image sets* $\operatorname{img}\mathrm{Cb}_n$ *are definable if and only if* $T_{\mathrm{Cb},0}$ *admits a model companion, if and only if this model companion is* T^b_{Cb} *(and is in fact a model completion).*

In case these equivalent conditions hold we shall simply denote the model companion by T_{Cb}, or, in the language $\mathcal{L}_P$, by T_P (although it is not usually a model companion in the language $\mathcal{L}_P$). In this case we say that the class of beautiful pairs of models of T is *weakly elementary*, in the sense that any sufficiently saturated model of the theory of this class, T_P, also belongs to it. By results of Poizat [Poi83], for a classical theory T this is further equivalent to T not having the finite cover property.

Corollary 4.6. *If T is $\aleph_0$-categorical then T^b_{Cb} is the model completion of $T_{\mathrm{Cb},0}$.*

Proof. In an $\aleph_0$-categorical theory every type-definable set is definable. $\blacksquare_{4.6}$

Since the theories discussed in previous sections (IHS, APr and $ALpL$) are $\aleph_0$-categorical, they satisfy the equivalent conditions of Theorem 4.4 in a somewhat uninteresting fashion. In the next section we consider a more interesting example of a non $\aleph_0$-categorical theory.

Just as we remarked in Section 1, these results can be extended to simple theories, where beautiful pairs are replaced with lovely pairs (see [BPV03]). The price to be paid is either to work with hyper-imaginary sorts (which can be done relatively smoothly in continuous logic) or to assume that uniform canonical bases exist in real or imaginary sorts (which is a strong form of elimination of hyper-imaginaries).

5. The case of $ACMVF$

A convenient feature of condition (iii) of Theorem 4.4 is that it remains invariant under the addition (or removal) of imaginary sorts. It may therefore serve as a criterion for the definability of the sets of uniform canonical bases even when these do not exist in any of the named sorts. As an example of this, let us consider the theory $ACMVF$ of algebraically closed metric valued fields, as defined in [Bena].

Theorem 5.1. *The equivalent conditions of Theorem 4.4 hold for $T = ACMVF$.*

Proof. We recall that a model of $ACMVF$ is not a valued field but rather a projective line $K\mathbf{P}^1$ over one, equipped with a modified distance function $d(x,y) = \|x-y\| = \frac{|x-y|}{\max|x|,1\cdot\max|y|,1}$ when $x, y \neq \infty$ and $d(x,\infty) = \frac{1}{\max|x|,1}$.

Given a polynomial $Q(Z,\bar W,\bar V)$ over $\mathbf{Z}$ and a tuple $\bar e$, let $\sqrt{Q(X,\bar e,P)}$ denote the collection of all roots of instances $Q(X,\bar e,\bar f)$ where $\bar f \in P$ (or more exactly, of the homogeneisation thereof, so ∞ can also be obtained a root). Let us define T_{Cb} to consist of $T_{\mathrm{Cb},0}$ along with the axioms saying that every $r \in [0,1]$, $b \in P$ and $Q(\bar e,\bar V)$ as above there is a such that $d(a,b) = d\big(a, \sqrt{Q(X,\bar e,P)}\big) = r$, or at least approximately so. One checks that $\sqrt{Q(X,\bar e,P)}$ is a definable set, so this is expressible in continuous logic.

In order to check that T_{Cb} is indeed a companion, let us consider an instance of the axioms. Since the map $x \mapsto x^{-1}$ is an isometric bijection of $K\mathbf{P}^1$, and for every Q there is Q' such that $\sqrt{Q(X,\bar e,P)}^{-1} = \sqrt{Q'(X,\bar e,P)}$, we may assume that $|b| \leq 1$. We may then add a such that $|a-c| = \max r, |b-c|$ for all $c \in \sqrt{Q(X,\bar e,P)} \smallsetminus \{\infty\}$, so in particular $|a| = \max|b|, r \leq 1$. It follows that $d\big(a, \sqrt{Q(X,\bar e,P)}\big) = d(a,b) = r$.

Now, let $(K\mathbf{P}^1, P)$ be an $\aleph_1$-saturated model of T_{Cb}, and we claim that the opposite holds, namely, that for all $b \in P \smallsetminus \{\infty\}$, countable set $\infty \notin A$ and $r \in (0,\infty)$, if B denotes the algebraic closure of $P \cup A$ then there exists a such that $|x-c| = \max r, |b-c|$ for all $c \in B$. We notice that if $|b| < r$ then replacing b with some $b' \in P$ such that $|b'| = r$, which exists by saturation, does not change the conditions of the problem. We may therefore assume that $|b| \geq r$. If $|b|, r \leq 1$ then (by saturation) there exists x such that $d(a,B) = d(a,b) = r$, and as in the previous paragraph this x is as desired. Otherwise, $|b| > 1$. We then find a such that for all $c \in B \smallsetminus \{\infty\}$, $|a-c| = \max r', |z-b^{-1}|$ with $r' = \frac{r}{|b|^2}$ (since $|b^{-1}|, r' \leq 1$). We observe that $|a| = \max|b^{-1}|, r|b^{-2}| = |b^{-1}|$, i.e., $|ab| = 1$. A direct calculation yields that for $c \in B \smallsetminus \{\infty\}$, $|a^{-1}-c| = \max r\left|\frac{c}{b}\right|, |b-c|$ (the case $c = 0$ may have to be considered separately). We now consider three cases (keeping in mind that $|b| \geq r$):

$$\max r \left|\frac{c}{b}\right|, |b-c| = \begin{cases} \max r \left|\frac{c}{b}\right|, |b| = |b| = \max r, |c-b|, & |c| < |b|, \\ \max r, |c-b|, & |c| = |b|, \\ \max r \left|\frac{c}{b}\right|, |c| = |c| = \max r, |c-b|, & |c| > |b|. \end{cases}$$

Our claim is thus proved.

Now let $p(x) \in \mathrm{S}_1(PA)$, and let q be any global extension of p to $\mathcal{M}$. Since $\mathcal{M}$ is $\aleph_1$-saturated, there exists $b \in M$ such that $r = d(x,b)^{q(x)}$ is minimal. Replacing p with $q{\restriction}_{PAb}$, we may assume that $b \in A$. By our previous claim, there exists $a \in M$ such that $d(a,b) = d(a,B) = r$, so this x must realise p.

We have thus shown that for every $\aleph_1$-saturated model $(\mathcal{M}, P)$ of the companion T_{Cb}, $\mathcal{M}$ is $\aleph_1$-saturated over P, so by Theorem 4.4 T_{Cb} is the model completion of $T_{\mathrm{Cb},0}$, and every uniform canonical base map has a definable image. $\blacksquare_{5.1}$

On the other hand, there are no canonical bases, so in particular no uniform ones, in the home sort of $ACMVF$. Indeed, we observed in [Bena] that 1-types over models are parametrised by spheres, and it is not difficult to see that if S is a sphere of non zero radius then in the home sort $\mathrm{dcl}(S) = \mathrm{dcl}(\varnothing)$. In the case of 1-types it is relatively easy to construct an imaginary sort in which uniform canonical bases exist. Indeed, let us consider the set of all pairs (a, r) with $a \in K\mathbf{P}^1$ and $r \in [0,1]$, more conveniently written as a_r, and let $[a_r]$ denote the closed ball of radius r around a. We define

$$d(a_r, b_s) = |r - s| \vee d(a,b) \dot{-} (r \wedge s).$$

Let us show that

$$d(a_r, c_t) \leq d(a_r, b_s) + d(b_s, c_t).$$

If $d(a_r, c_t) = |r - t|$ then the inequality is clear. Otherwise, may assume that $d(a,c) \leq d(a,b)$, so

$$d(a_r, c_t) = d(a,c) \dot{-} (r \wedge t) \leq d(a,b) \dot{-} (r \wedge s) + |s - t| \leq d(a_r, b_s) + d(b_s, c_t).$$

(A somewhat less direct way of observing the same would consist of defining $\varphi(x,a,r) = d(x,a) \dot{-} r$ (namely $d(x,[a_r])$), and observe that $d(a_r, b_s) = \sup_x |\varphi(x,a,r) - \varphi(x,b,s)|$. Clearly $d(a_r, b_s) = 0$ if and only if $[a_r] = [b_s]$.) On the other hand, the set of all such quotients is incomplete, and the set of completions consists, in addition to all closed balls, of all spheres over the structure (this construction can be carried out in any bounded ultra-metric structure). In this imaginary sort, 1-types admit a uniform canonical base map. The case of canonical bases for n-types and general elimination of imaginaries for $ACMVF$ are far more complex, compare with [HHM06] as well as with more recent results of Hrushovski and Loeser with respect to uniform canonical bases of generically stable types in $ACVF$.

References

[BBHa] Itaï Ben Yaacov, Alexander Berenstein, and C. Ward Henson, *Almost indiscernible sequences and convergence of canonical bases*, submitted, arXiv:0907.4508.

[BBHb] ———, *Model-theoretic independence in the Banach lattices $L_p(\mu)$*, Israel Journal of Mathematics, to appear, arXiv:0907.5273.

[Bena] Itaï Ben Yaacov, *Model theoretic properties of metric valued fields*, submitted, arXiv:0907.4560.

[Benb] ———, *On theories of random variables*, submitted, arXiv:0901.1584.

[Ben06] ———, *Schrödinger's cat*, Israel Journal of Mathematics **153** (2006), 157–191, doi:10.1007/BF02771782.

[BPV03] Itaï Ben Yaacov, Anand Pillay, and Evgueni Vassiliev, *Lovely pairs of models*, Annals of Pure and Aplied Logic **122** (2003), 235–261, doi:10.1016/S0168-0072(03)00018-6.

[BU] Itaï Ben Yaacov and Alexander Usvyatsov, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society, to appear, arXiv:0801.4303.

[HHM06] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson, *Definable sets in algebraically closed valued fields: elimination of imaginaries*, Journal für die Reine und Angewandte Mathematik. [Crelle's Journal] **597** (2006), 175–236, doi:10.1515/CRELLE.2006.066.

[HKP00] Bradd Hart, Byunghan Kim, and Anand Pillay, *Coordinatisation and canonical bases in simple theories*, Journal of Symbolic Logic **65** (2000), 293–309, doi:10.2307/2586538.

[Hru97] Ehud Hrushovski, *Simplicity and the Lascar group*, unpublished, 1997.

[LT79] Joram Lindenstrauss and Lior Tzafriri, *Classical Banach spaces. II*, Ergebnisse der Mathematik und ihrer Grenzgebiete [Results in Mathematics and Related Areas], vol. 97, Springer-Verlag, Berlin, 1979, Function spaces.

[Pil96] Anand Pillay, *Geometric stability theory*, Oxford Logic Guides, vol. 32, The Clarendon Press Oxford University Press, New York, 1996, Oxford Science Publications.

[Poi83] Bruno Poizat, *Paires de structures stables*, Journal of Symbolic Logic **48** (1983), no. 2, 239–249, doi:10.2307/2273543.

[She90] Saharon Shelah, *Classification theory and the number of nonisomorphic models*, second ed., Studies in Logic and the Foundations of Mathematics, vol. 92, North-Holland Publishing Co., Amsterdam, 1990.

Itaï Ben Yaacov, Université Claude Bernard – Lyon 1, Institut Camille Jordan, CNRS UMR 5208, 43 boulevard du 11 novembre 1918, 69622 Villeurbanne Cedex, France

URL: http://math.univ-lyon1.fr/~begnac/