

ON UNIFORM CANONICAL BASES IN L_p LATTICES AND OTHER METRIC STRUCTURES

ITAI 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 be 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 [BBH], 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.

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 [BU10].

2000 *Mathematics Subject Classification.* 03C45, 46B42, 12J25.

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

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

Revision –revision– of 21st May 2012.

1. UNIFORM CANONICAL BASES

In classical logic, stable theories are characterised by the property that for every model \mathcal{M} , every type $p(\bar{x}) \in S_{\bar{x}}(\mathcal{M})$ 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 $\bar{b} \in M$:

$$\varphi(\bar{x}, \bar{b}) \in p \quad \iff \quad \models \psi(\bar{b}).$$

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

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

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 $\text{Cb}_{\varphi}(p)$. 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 $\text{Cb}(p)$. 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

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

We shall hereafter refer to definable predicates as formulae as well, since for our present 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 [BBH11] (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 $p(\bar{x}) \in S_{\bar{x}}(\mathcal{M})$,

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

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 $\text{Cb}(p)$ to be merely a set which is only known up to inter-definability, as is the common practice, then the expression $d\varphi(\bar{y}, \text{Cb}(p))$ 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 $\mathcal{M} \models T$ there exists a tuple $A \subseteq M^{eq}$ in the sort of Z such that for each $\varphi(\bar{x}, \bar{y})$:

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

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

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 $\text{Cb}(p)$ 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 $\text{Cb}(p) = B$ 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 [BU10]. $\blacksquare_{1.2}$

Lemma 1.3. *The image img 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 img 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 $\text{Cb}' = f \circ \text{Cb}$ 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 $d\varphi(\bar{y}, Z)$ is a formula then $d\varphi(\bar{y}, f^{-1}(W))$ 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 $f: \text{Cb}(p) \mapsto \text{Cb}'(p)$ 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.

Notation 1.5. When \mathcal{M} is a model and \bar{a} a tuple in some elementary extension, we write $\text{Cb}(\bar{a}/M)$ for $\text{Cb}(p)$ where $p = \text{tp}(\bar{a}/M)$.

Lemma 1.6. *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 $f^{\text{Cb}}(\text{Cb}(\bar{a}/M), \bar{b}) = \text{Cb}(f(\bar{a}, \bar{b})/M)$ is definable as well for $(\bar{a}, \bar{b}) \in \text{dom } f$, $\bar{b} \in M$, 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 $\text{tp}(f(\bar{a}, \bar{b})/M)$ by

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

where $\psi(\bar{x}, \bar{y}\bar{z}) = \varphi(f(\bar{x}, \bar{y}), \bar{z})$. The case with parameters follows. $\blacksquare_{1.6}$

Lemma 1.7. *Let Cb be a uniform canonical base map, say on the sort of n -tuples, into some infinite sort, and let $\text{Cb}(p)_i$ denote its i th coordinate. Then the map $\bar{a} \mapsto \text{Cb}(\bar{a}/M)_i$ is uniformly continuous, and uniformly so regardless of \mathcal{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. The general case follows using Lemma 1.4 and the fact that definable functions are uniformly continuous. $\blacksquare_{1.7}$

Remark 1.8. 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 1.9. 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 1.10. 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 $p = \text{tp}(\bar{v}/E)$ is the orthogonal projection $P_E(\bar{v})$. Indeed, by quantifier elimination it will be enough to show that $p(\bar{x})$ admits a φ -definition over $P_E(\bar{v})$ for every

formula of the form $\varphi(\bar{x}, y) = \|\sum \lambda_i x_i + y\|^2$, for any choice of scalars $\bar{\lambda}$. We then observe that for all $u \in E$ we have

$$\varphi(\bar{v}, u) = \left\| \sum \lambda_i v_i \right\|^2 - \left\| \sum \lambda_i P_E(v_i) \right\|^2 + \left\| \sum \lambda_i P_E(v_i) + u \right\|^2.$$

The second and third terms are definable from u using $P_E(\bar{v})$ as parameters, while the first term is simply a constant which does not depend on u .

Since $P_E(\bar{v})$ alone does not allow us to recover $\|\sum \lambda_i v_i\|^2$, and therefore does not allow us to recover p either, this canonical base is not uniform. By adding the missing information (namely, the inner product on the v_i , in the sort $[-1, 1]$) we obtain a uniform canonical base:

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

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 1.11. 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 $\text{tp}(\bar{v}/E) \mapsto P_E(\bar{v})$ is weakly uniform.

Proposition 1.12. *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 $f \circ \text{Cb}'$, where Cb' is a uniform canonical base map with target sort W and $f: \text{img } \text{Cb}' \rightarrow Z_0$ is definable. Let Φ be the set of all formulae $\varphi(\bar{x}, W)$. For every such formula, the value $\varphi(\bar{a}, \text{Cb}'(\bar{a}/M))$ is uniformly definable from $\text{Cb}'(\bar{a}/M)$, call it $g_\varphi(\text{Cb}'(\bar{a}/M))$, and let $g = (g_\varphi)_{\varphi \in \Phi}$. Then $(f, g): \text{img } \text{Cb}' \rightarrow Z_0 \times Z_1$ 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 $\text{Cb} = (f, g) \circ \text{Cb}'$ 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 $C = \text{Cb}'(\bar{a}/M)$, $D = \text{Cb}'(\bar{b}/M)$, 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. $\blacksquare_{1.12}$

Thus our question can be restated as

Question 1.13. 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 1.12 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. There is a caveat, though, namely a uniform canonical base exists not in the home sort of events (see Corollary 2.5), but in the (still mathematically natural) imaginary sort of $[0, 1]$ -valued random variables.

In [Benb] we define the theory *ARV* of atomless spaces of $[0, 1]$ -valued random variables, in the language $\{0, \neg, \frac{1}{2}, \vee, \wedge\}$, where 0 is a constant, $\neg X = 1 - X$ and $\frac{1}{2}X$ are unary, the lattice operations \vee and \wedge are binary, and integration is recovered as $\mathbf{E}[X] = d(X, 0)$. We show there that *ARV* is bi-interpretable with *APr*, the theory of atomless probability algebra: to a probability algebra \mathcal{A} one associates the space $\mathcal{M} = L^1(\mathcal{A}, [0, 1])$ of $[0, 1]$ -valued \mathcal{A} -measurable random variables, from which one can recover $\mathcal{A} = \sigma(M)$. In addition, *ARV* is \aleph_0 -stable, eliminates quantifiers, and admits definable continuous calculus: if $\tau: [0, 1]^n \rightarrow [0, 1]$ is any continuous function then the map $\bar{X} \mapsto \tau(\bar{X})$ is definable.

Thus, in order to construct uniform canonical bases for APr in the sort of $[0, 1]$ -valued random variables, it will be enough (and in a sense, better) to show that ARV admits uniform canonical bases in the home sort.

For models $\mathcal{M} \preceq \mathcal{N} \models ARV$ and $X \in N$ let us write $\mathbf{E}[X|M]$ for the conditional expectation $\mathbf{E}[X|\sigma(M)]$, itself a member of M . For tuples $\bar{X} \in N^n$ and $\bar{k} \in \mathbf{N}^n$, let $\bar{X}^{\bar{k}} = \prod X_i^{k_i}$.

Lemma 2.1. *Let $\mathcal{M} \preceq \mathcal{N} \models ARV$ 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 $\text{Cb}(\bar{X}/M)$. Also, for every $\bar{\ell} \in \mathbf{N}^m$, $\mathbf{E}[\bar{X}^{\bar{k}}\bar{Y}^{\bar{\ell}}]$ is uniformly definable from $\mathbf{E}[\bar{X}^{\bar{k}}|M]$ and $\bar{Y} \in M^m$.*

Proof. By the definable continuous calculus, the function $(\bar{X}, Y) \mapsto |\bar{X}^{\bar{k}} - Y|^2$ is uniformly definable, and by Lemma 1.6 the predicate $\|\bar{X}^{\bar{k}} - y\|_2$ is uniformly definable for $y \in M$ from $\text{Cb}(\bar{X}/M)$. For $Y \in M$ we have

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

where the infimum is taken in M . Thus the graph of the function $\text{Cb}(\bar{X}/M) \mapsto \mathbf{E}[\bar{X}^{\bar{k}}|M]$ is type-definable in \mathcal{M} , whence it follows that the function itself is definable, and uniformly so in all models of ARV . For the second assertion, use definable continuous calculus and the fact that for $\bar{Y} \in M^m$ we have $\mathbf{E}[\bar{X}^{\bar{k}}\bar{Y}^{\bar{\ell}}] = \mathbf{E}[\mathbf{E}[\bar{X}^{\bar{k}}|M] \cdot \bar{Y}^{\bar{\ell}}]$. $\blacksquare_{2.1}$

Theorem 2.2. *For n -types over models in ARV ,*

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

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

Proof. By Lemma 2.1, 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 Lemma 2.1 again, this tuple determines $\mathbf{E}[\bar{X}^{\bar{k}}\bar{Y}^{\bar{\ell}}]$ for every $\bar{Y} \in M^m$ and $\bar{\ell} \in \mathbf{N}^m$. By the Stone-Weierstraß Theorem, every continuous function $\tau: [0, 1]^{n+m} \rightarrow [0, 1]$ is uniformly approximated by polynomials, so $\mathbf{E}[\tau(\bar{X}, \bar{Y})]$ is determined for every such τ and \bar{Y} , which, by quantifier elimination, is enough to determine $\text{tp}(\bar{X}/M)$. $\blacksquare_{2.2}$

Corollary 2.3. *For n -types over models in APr ,*

$$\text{Cb}(\bar{A}/\mathcal{A}) = \left(\mathbf{P} \left[\bigwedge_{i \in s} A_i \mid \mathcal{A} \right] \right)_{\emptyset \neq s \subseteq n}$$

is a uniform canonical base (in the sort of $[0, 1]$ -valued random variables), where $\mathbf{P}[\cdot|\mathcal{A}]$ denotes conditional probability.

Proof. Immediate. $\blacksquare_{2.3}$

Let us now show that this cannot be improved.

Lemma 2.4. *There is no \emptyset -definable injective map from a model $\mathcal{M} \models ARV$ to $\sigma(M)^\kappa \times [0, 1]^\lambda$ (for any cardinals κ, λ).*

Proof. Assume there were one, say $(f, g): M \rightarrow \sigma(M)^\kappa \times [0, 1]^\lambda$, where f goes to $\sigma(M)^\kappa$ and g to $[0, 1]^\lambda$. Let $A \in \sigma(M)$ have measure one half, with indicator function $\mathbf{1}_A$, and for $0 \leq r \leq 1$ let $X_r = r\mathbf{1}_A \in M$. For $i < \kappa$, the event $f_i(X_r)$ must be definable from X_r , namely from A , and therefore must belong to the discrete set $\{0, 1, A, \neg A\}$. Since the map f_i , being definable, must be continuous, we conclude that $r \mapsto f_i(X_r)$ must be constant, so $f(\mathbf{1}_A) = f(0)$.

Now let $B \in \sigma(M)$ be any other event of measure one half. Then $f(\mathbf{1}_A) = f(0) = f(\mathbf{1}_B)$, and by quantifier elimination $\mathbf{1}_A \equiv \mathbf{1}_B$, whereby $g(\mathbf{1}_A) = g(\mathbf{1}_B)$, so (f, g) is not injective after all. $\blacksquare_{2.4}$

Corollary 2.5. *The theory APr does not admit a weakly uniform canonical base in its home sort (of events).*

Proof. Let $\mathcal{A} \models APr$, say the probability algebra of Ω , and let \mathcal{B} be the probability algebra of $\Omega \times [0, 1]$. Then $\mathcal{B} \models APr$ as well, and identifying $A \in \mathcal{A}$ with $A \times [0, 1]$, we obtain an embedding $\mathcal{A} \subseteq \mathcal{B}$ which is elementary by quantifier elimination of APr . More generally $X \in L^1(\mathcal{A}, [0, 1])$, let $B_X = \{(\omega, r) : r \leq X(\omega)\} \in \mathcal{B}$ (which is well defined, up to null measure). Thus, with Cb a uniform canonical base map as in Corollary 2.3 (for types of singletons), we have $\text{Cb}(B_X/\mathcal{A}) = \mathbf{P}[B_X|\mathcal{A}] = X$, so $\text{img Cb} = L^1(\mathcal{A}, [0, 1])$.

Now, assume there existed a weakly uniform canonical base map for types of singletons in APr with values in the home sort. Then by Proposition 1.12 there would exist a uniform canonical base map

with values in the home sort and $[0, 1]$, and by Lemma 1.4 it is of the form $f \circ \text{Cb}$, where $f: \text{img Cb} \rightarrow \mathcal{A}^\kappa \times [0, 1]^\lambda$ is definable and injective. This (together with the bi-interpretability of \mathcal{A} with $L^1(\mathcal{A}, [0, 1])$) contradicts Lemma 2.4. $\blacksquare_{2.5}$

3. UNIFORM CANONICAL BASES IN ATOMLESS L_p LATTICES

Recall that L_pL denotes the theory of L_p lattices for some fixed $p \in [1, \infty)$, and that AL_pL denotes the theory of atomless ones: their models are exactly those abstract Banach lattices which are isomorphic to a Banach lattice of the form $L_p(\Omega)$, where Ω is a measure space, atomless in the case of AL_pL . We may write the measure space more explicitly as (Ω, Σ, μ) , where Ω is a set of points, Σ a σ -algebra on Ω and μ a σ -additive measure on Σ , atomless in the case of AL_pL .

Stability, independence and related notions were studied for AL_pL by Berenstein, Henson and the author in [BBH11]. Here we shall consider Banach lattices (as we do all Banach space structures) via their unit balls (from which the entire Banach lattice can be recovered), in the language $\{0, -, \frac{x+y}{2}, |\cdot|\}$, where, for æsthetic reasons, we halve the distance: $d(x, y) = \|\frac{x-y}{2}\|$ (see [Ben09]). The results of [BBH11] were given in a different, yet equivalent, formalism, so they hold fully in ours. The theory AL_pL 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 $t: \mathbf{R}^n \rightarrow \mathbf{R}$ 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 $\varphi: \mathbf{R}^n \rightarrow \mathbf{R}$ 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

$$t(\bar{z}) = \frac{bx_1 - ay_1}{x_1 y_0 - x_0 y_1} z_0 + \frac{ay_0 - bx_0}{x_1 y_0 - x_0 y_1} z_1.$$

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 $t_k \rightarrow \varphi$ 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 $\varphi: E^n \rightarrow E$ is uniformly definable across all (unit balls of) 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 L_pL and L_qL are quantifier-free bi-definable. 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. The map $\theta: f \mapsto f^{\frac{p}{q}}$ clearly is a bijection between $L_p(\Omega)$ and $L_q(\Omega)$. For any definable map or predicate in $L_q(\Omega)$, call it τ , let τ^{L_q} denote its pre-image in $L_p(\Omega)$ by θ . Thus, for example, $\|x\|^{L_q} = \|x\|^{\frac{p}{q}}$, so θ is also a bijection between the respective unit balls. Also, $0^{L_q} = 0$, and similarly for $-$, and $|\cdot|$ (namely, they commute with θ). For $\frac{x+y}{2}$, we have $(\frac{x+y}{2})^{L_q} = \left(\frac{x^{p/q} + y^{p/q}}{2}\right)^{\frac{q}{p}}$: as a map $\mathbf{R}^2 \rightarrow \mathbf{R}$,

this is continuous and \mathbf{R}^+ -homogeneous of degree one, and therefore uniformly definable in all Banach lattices, and in particular in models of L_pL . Finally, $d(x, y)^{L_q} = \|\frac{x-y}{2}\|^{L_q} = \|(\frac{x-y}{2})^{L_q}\|^{L_q}$, and we already know that $(\frac{x-y}{2})^{L_q}$ is definable. It follows that the predicate $f \mapsto \int f^p$ is definable in $L_p(\Omega)$: when $p = 1$ this is merely $\|f^+\| - \|f^-\|$, where $f^+ = \frac{|f|+f}{2}$ and $f^- = \frac{|f|-f}{2}$, and for the general case we use the bi-definability proved above with $q = 1$. By the same arguments, $(x, y) \mapsto x^{\frac{1}{q}}y^{\frac{1}{q'}}$ is definable in $L_p(\Omega)$. Given $f, g \in L_p(\Omega)$ which represent $f^{\frac{p}{q}} \in L_q(\Omega)$ and $g^{\frac{p}{q'}} \in L_{q'}(\Omega)$, respectively, we have $\langle f^{\frac{p}{q}}, g^{\frac{p}{q'}} \rangle = \left(\int f^{\frac{1}{q}}g^{\frac{1}{q'}} \right)^p$, which is definable from $f^{\frac{1}{q}}g^{\frac{1}{q'}}$ and therefore from (f, g) , concluding the proof. $\blacksquare_{3.3}$

Recall from [BBH11] that for any embedding $E \subseteq F$ of L_p lattices there are presentations of E and F as $L_p(\Omega_1, \Sigma_1, \mu_1) \subseteq L_p(\Omega_2, \Sigma_2, \mu_2)$, such that $\Omega_1 \subseteq \Omega_2$, $\Sigma_1 \subseteq \Sigma_2$ and μ_2 extends μ_1 . For $g \in E$, $f \in F$, and presentation as above, we say that $g = \mathbf{E}[f|\Sigma_1]$ if $\int_A g = \int_A f$ for all $A \in \Sigma_1$ of finite measure. Uniqueness of g is clear, and existence is fairly standard (at least under some mild hypotheses on the measure spaces in question). It is shown in [BBH11] that, moreover, $\mathbf{E}[f|\Sigma_1]$ does not depend on the choice of presentation (i.e., of measure spaces), and is therefore denoted $g = \mathbf{E}[f|E]$, the *conditional expectation* of f with respect to E , which will be an essential tool for this section. We shall see below (Lemma 3.6) that for $p > 1$, the Krivine calculus provides us with a more elementary argument showing that the conditional expectation depends only on $\text{tp}(f/E)$ (and not on the presentation), and moreover definably so (an even more elementary argument shows independence of the presentation when $p = 1$, see Remark 3.7).

Let us recall from [BBH11] a few facts regarding types in AL_pL . Recall first that f, g in a Banach lattice are called *orthogonal*, in symbols $f \perp g$, if $|f| \wedge |g| = 0$, and if $E \subseteq F$ are two Banach lattices then $E^\perp = \{f \in F : f \perp E\} = \{f \in F : f \perp g \text{ for all } g \in E\}$ is a sub-lattice of F . Recall also that if $E \subseteq F$ are L_p lattices then $F = E^{\perp\perp} \oplus E^\perp$, so each $f \in F$ can be written uniquely as $f = f|_E + f|_{E^\perp}$, where $f|_E \perp E$ and $f|_{E^\perp} \perp E^\perp$.

Fact 3.4. *By [BBH11, Lemma 3.8], $\text{tp}(f/E)$ is determined by $\text{tp}(f|_E/E)$ together with $\|f|_{E^\perp}^\pm\|$, or equivalently, together with $\|f^\pm\|$. If $E = L_p(\Omega_1, \Sigma_1, \mu_1) \subseteq F = L_p(\Omega_2, \Sigma_2, \mu_2)$ is a concrete representation as above then $f|_E = f\mathbf{1}_{\Omega_1}$ and $f|_{E^\perp} = f\mathbf{1}_{\Omega_2 \setminus \Omega_1}$, and by [BBH11, Proposition 3.7], $\text{tp}(f|_E/E)$ is determined by the Σ_1 -conditional distribution of $f|_E$, namely by the map $t \mapsto \mathbf{P}[f > t|\Sigma_1]$.*

We say that F is *atomless over E* if for every $g \in E$, $g \geq 0$, there is $f \in F$, $0 \leq f \leq g$, such that $f \perp g - f$ and $\|f\| = \|g\|/2$, or equivalently, if for any concrete presentation as above, μ_2 is atomless over Σ_1 . If this is the case then, by virtue of the above, together with [BBH11, Fact 2.15], for any f in an extension of E , $\text{tp}(f|_E/E)$ is realised in F . If, in addition, $E^\perp \subseteq F$ is non trivial, then any 1-type (or indeed, n -type, but we shall not require this) over E is realised in F .

Let Ω be a measure space, and let $\mathcal{A}(\Omega)$ be the algebra (or more precisely, the relatively complemented distributive lattice) of finite measure sets modulo null measure sets, which, with some abuse of terminology, we shall call the *measure algebra* of Ω . Let Ω' be the Stone space of $\mathcal{A}(\Omega)$, an extremally disconnected locally compact space, equipped with the unique regular Borel measure which associates to each compact open set the (finite) measure of the corresponding measurable subset of Ω , which exists by Carathéodory's Theorem.

Fact 3.5. *With Ω and Ω' as above, $\mathcal{A}(\Omega')$ is canonically isomorphic to $\mathcal{A}(\Omega)$, yielding a canonical identification $L_p(\Omega) \cong L_p(\Omega')$, so we may assume that $\Omega = \Omega'$. Such a space has the particular property that every $g \in L_p(\Omega)$ has a unique continuous representative, so it makes sense to speak of the value $g(\omega) \in [-\infty, \infty]$ at some $\omega \in \Omega$, which is finite for almost all ω .*

Algebraic and lattice operations, and therefore the lattice order, are reflected in these continuous representatives. If $A \subseteq L_p(\Omega)$ is countable (or separable) then A is bounded in $L_p(\Omega)$ if and only if the point-wise supremum $g(\omega) = \sup_{f \in A} f(\omega)$ is in L_p , in which case it is equal almost everywhere to the continuous representative of $\sup A$.

Proof. The construction is straightforward. To see that it has the desired properties, it is enough to consider restrictions to finite measure sub-spaces, in which case this is shown in [Ben06, Theorem 1.9 and Proposition 1.14]. For a more elaborate discussion, see [Fre04, Chapter 32]. $\blacksquare_{3.5}$

This spares us the headache of choosing representatives: for example, given a convex function $x \mapsto \Psi_f(x) \in L_p(\Omega)$, it is automatically true that $x \mapsto \Psi_f(x)(\omega)$ is convex for each ω .

Let us fix $E \models ALpL$ and a type $q = \text{tp}(f/E)$ for some f in an extension of E . Let us also fix a concrete representation $E = L_p(\Omega)$, with Ω as in Fact 3.5, and define $E' = L_p((\Omega \cup \{+, -\}) \times [0, 1])$, where the interval is taken with the Lebesgue measure, and each of \pm has measure one. Let $\pi_0: \Omega \times [0, 1] \rightarrow \Omega$ denote the projection, and identify a function $g: \Omega \rightarrow \mathbf{R}$ with $g \circ \pi_0: \Omega \times [0, 1] \rightarrow \mathbf{R}$, as well as with its extension by zeroes to $(\Omega \cup \{\pm\}) \times [0, 1]$, yielding a canonical embedding $E \subseteq E'$, which, by quantifier elimination for $ALpL$, is elementary. By the previous paragraph, we may assume that $f \in E'$. The conditional expectation of f with respect to Ω (namely, to its Borel algebra) is

$$\mathbf{E}[f|\Omega](\omega) = \int_0^1 f(\omega, t) dt,$$

and $\|\mathbf{E}[f|\Omega]\| \leq \|f\|$ by Jensen's inequality and convexity of $x \mapsto x^p$ (for $f \geq 0$ the integral is finite almost everywhere, by Fubini's Theorem and Jensen's inequality, so it is defined almost everywhere for all f).

Lemma 3.6. *Assume that $p > 1$, and let $E \models ALpL$, $q \in S_1(E)$, and $f \models q$ in E' as above. Then the conditional expectation $\mathbf{E}[f|\Omega]$ defined above depends only on q , that is to say that for any presentation $E = L_p(\Omega') \subseteq L_p(\Omega'') = F$, and any $f' \in F$, if $f' \models q$ then $\mathbf{E}[f'|\Omega'] = \mathbf{E}[f|\Omega]$. Moreover, for any uniform canonical base map Cb , the map $\text{Cb}(f/E) \mapsto \mathbf{E}[f|E]$ (where $\mathbf{E}[f|E]$ now denotes $\mathbf{E}[f|\Omega]$) is definable in $ALpL$. More generally, if $t(x, \bar{y})$ is any lattice term, then the map $(\text{Cb}(f/E), \bar{g}) \mapsto \mathbf{E}[t(f, \bar{g})|E]$, where $\bar{g} \in E$, is definable.*

Proof. Clearly, if $g \in E$ satisfies that $\langle f, h \rangle = \langle g, h \rangle$ for all $h \in E^* = L_{p'}(\Omega)$, where $p' = \frac{p}{p-1}$ then $g = \mathbf{E}[f|\Omega]$, and it is fairly standard to check that the converse holds as well. It follows that $g = \mathbf{E}[f|\Omega]$ if and only if $\int fh^{p-1} = \int gh^{p-1}$ for all $h \in E$, or equivalently, for all h in the unit ball of E . The map $(x, y) \mapsto x^{\frac{1}{p}}y^{\frac{p-1}{p}}$ is uniformly definable by the Krivine calculus, so $(x, y) \mapsto \|x^{\frac{1}{p}}y^{\frac{p-1}{p}}\|^p = \int |x||y|^{p-1}$ is a definable predicate. Separating into positive and negative parts we see that $\varphi(x, y) = \int xy^{p-1}$ is a definable predicate as well. Together with our characterisation of $\mathbf{E}[f|\Omega]$, this implies that it indeed only depends on $\text{tp}(f/E)$ (in fact, all of this holds even when E is a mere model of LpL , not necessarily atomless, in which case $\mathbf{E}[f|\Omega]$ only depends on the quantifier-free type of f over E).

From now on we write $\mathbf{E}[f|E]$ for $\mathbf{E}[f|\Omega]$. With φ as above, in E we have $x = \mathbf{E}[f|E]$ if and only if

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

and this remains true if we restrict z to the unit ball. Therefore the graph of $\text{Cb}(f/E) \mapsto \mathbf{E}[f|E]$ is type-definable, and moreover, given a bound on $\|f\|$ we have a bound on $\|\mathbf{E}[f|E]\|$ (so, in a sense, we know in which ‘‘sort’’ $\mathbf{E}[f|E]$ lives), so the map is definable. The case of a lattice term follows by Lemma 1.6. ■_{3.6}

Remark 3.7. This argument does not work for $p = 1$, and we shall see below that for $p = 1$ the definability of $\text{Cb}(f/E) \mapsto \mathbf{E}[f|E]$ fails. However, when $p = 1$ we have $\int f = \|f^+\| - \|f^-\|$, and if $h \in E$ is positive, say with support A in some presentation, then $\int f \upharpoonright_A = \lim_{N \rightarrow \infty} (f \wedge Nh) \vee (-Nh)$, so again $\mathbf{E}[f|E]$ only depends on $\text{tp}(f/E)$.

Given f in some extension of E , let $f_0 = \mathbf{E}[|f||E] \in E$, and consider the function $\Psi_f: \mathbf{R} \rightarrow E$, $\Psi_f(x) = \mathbf{E}[xf_0 \dot{-} f|E]$. It is convex, and therefore admits one-sided derivatives

$$D_x^+ \Psi_f = \lim_{\varepsilon \searrow 0} \frac{\Psi_f(x + \varepsilon) - \Psi_f(x)}{\varepsilon} = \inf_{\varepsilon > 0} \frac{\Psi_f(x + \varepsilon) - \Psi_f(x)}{\varepsilon},$$

and similarly

$$D_x^- \Psi_f = \sup_{\varepsilon > 0} \frac{\Psi_f(x - \varepsilon) - \Psi_f(x)}{-\varepsilon}.$$

The function Ψ_f is determined by $\text{tp}(f/E)$, and by Lemma 3.6, for $p > 1$ it is even uniformly definable from the canonical base. Conversely,

Convention 3.8. Given f in an extension of E and $f_0 = \mathbf{E}[|f||E]$ as above, one can always choose a concrete representation $E = L_p(\Omega)$ such that $f_0 = \mathbf{1}_A$ is an indicator function. We may moreover assume that Ω is as in Fact 3.5, in which case A is a compact open set. We fix such a presentation for the rest of the section.

We may still assume that $f \in E' = L_p((\Omega \cup \{+, -\}) \times [0, 1])$, as in the discussion preceding Lemma 3.6.

Lemma 3.9. *Under Convention 3.8, the conditional distribution of $f|_E$ with respect to Ω is determined by*

$$\mathbf{P}[f|_E \leq x|\Omega] \cdot \mathbf{1}_A = \mathbf{P}[f \leq xf_0|\Omega] \cdot f_0 = D_x^+ \Psi_f, \quad f|_E \cdot (1 - \mathbf{1}_A) = 0,$$

In particular, the triplet $(\Psi_f, \|f^\pm\|)$, determines $\text{tp}(f/E)$.

Proof. Indeed, for $\omega \in A$ we have

$$\Psi_f(x)(\omega) = \mu_{[0,1] \times \mathbf{R}} \left\{ (t, s) : f(\omega, t) \leq s \leq x \right\}$$

and

$$\begin{aligned} D_x^+ \Psi_f(\omega) &= \inf_{\varepsilon} \frac{1}{\varepsilon} \mu_{[0,1] \times \mathbf{R}} \left\{ (t, s) : f(\omega, t) \vee x \leq s \leq x + \varepsilon \right\} \\ &= \mu_{[0,1]} \left\{ t : f(\omega, t) \leq x \right\} = \mathbf{P}[f|_E \leq x|\Omega](\omega). \end{aligned}$$

The second assertion now follows from Fact 3.4. ■_{3.9}

A useful tool in the study of convex functions is the *Legendre transform*, which allows to pass between the spaces of points and of slopes. For a function $\varphi: \mathbf{R} \rightarrow \mathbf{R} \cup \{\infty\}$, not identically ∞ , we define

$$\varphi^*(t) = \sup_{x \in C_\varphi} tx - \varphi(x) \in \mathbf{R} \cup \{\infty\}, \quad \text{where } C_\varphi = \{x \in \mathbf{R} : \varphi(x) < \infty\}.$$

The function φ^* is always convex and continuous on $\overline{C_{\varphi^*}}$ (namely at the boundary, since, being convex, it is automatically continuous in the interior). If φ^* is not identically ∞ , which is always the case when φ is convex, we can define the double Legendre transform φ^{**} , which is the greatest function below φ which is convex and continuous on $\overline{C_\varphi}$. If φ is convex and continuous on $\overline{C_\varphi}$ then $\varphi = \varphi^{**}$. In this case, we have for every $x \in C_\varphi$ and $t \in C_{\varphi^*}$:

$$(1) \quad \varphi(x) + \varphi^*(t) = tx \iff D_t^- \varphi^* \leq x \leq D_t^+ \varphi^* \iff D_x^- \varphi \leq t \leq D_x^+ \varphi.$$

(At the boundary of the domain, the derivatives may be $\pm\infty$, and are always finite in the interior.) This can be generalised to functions on spaces of higher dimension, with tx replaced with a duality pairing, see for example in Rockafellar [Roc70].

Working with a presentation $E = L_p(\Omega)$ as per Convention 3.8, we can deduce similar properties for the lattice-valued Ψ_f .

Lemma 3.10. *Let $f \in F \supseteq E$, $f_0 = \mathbf{E}[|f||E]$ and $\Psi_f(x) = \mathbf{E}[xf_0 \div f]$ be as above, and define*

$$\Psi_f^*(t) = \sup_x txf_0 - \Psi_f(x), \quad \Psi_f^{**}(x) = \sup_{t \in [0,1]} txf_0 - \Psi_f^*(t).$$

Then Ψ_f^ is finite, convex and continuous on $[0, 1]$, with $\Psi_f^*(0) = 0$ and $\Psi_f^*(1) = \mathbf{E}[f|E]$, it is infinite for $t \notin [0, 1]$, and $\Psi_f = \Psi_f^{**}$.*

Proof. We have $\|\Psi_f(x)\|_\infty \leq |x| + 1$ for $x \in \mathbf{R}$, so for $\omega \in \Omega$ we may define a convex function $\psi_\omega: \mathbf{R} \rightarrow \mathbf{R}$ by $\psi_\omega(x) = \Psi_f(x)(\omega)$. For $x \leq y$ we have $\Psi_f(x) \leq \Psi_f(y) \leq \Psi_f(x) + (y - x)f_0$, and it follows that $\psi_\omega(x) \leq \psi_\omega(y) \leq \psi_\omega(x) + (y - x)$. Passing to the Legendre transform, this means that $\Psi_f^*(t) = \infty$ and $\psi_\omega^*(t) = \infty$ for all ω and $t \notin [0, 1]$. For $t = 0, 1$ we have

$$\Psi_f^*(0) = \lim_{x \rightarrow -\infty} -\mathbf{E}[xf_0 \div f|E] = 0, \quad \Psi_f^*(1) = \lim_{x \rightarrow \infty} \mathbf{E}[xf_0 \wedge f|E] = \mathbf{E}[f|E].$$

By convexity, $\Psi_f^*(t) \in E$ (i.e., is finite) for each $t \in [0, 1]$, and by Fact 3.5, $\psi_\omega^*(t) = \Psi_f^*(t)(\omega)$ for almost all ω . Using the continuity of convex functions, we can exchange the quantifiers: for almost all ω , $\psi_\omega^* = \Psi_f^*(\cdot)(\omega)$. It follows that for each x , for almost all ω : $\psi_\omega^{**}(x) = \Psi_f^{**}(x)(\omega)$. Thus $\Psi_f(x) = \Psi_f^{**}(x)$, as desired.

Let us show that $\Psi_f^*(t) \rightarrow 0$ as $t \searrow 0$ in L_p (point-wise convergence almost everywhere is already known). For the positive part this just follows from convexity. For the negative part, for every $\varepsilon > 0$ there exists $s > 0$ such that $\|(-sf_0) \div f\| < \varepsilon \|f_0\|$, so for $0 < \delta < \varepsilon/s$ we have

$$\Psi_f^*(\delta) \geq -\delta sf_0 - \mathbf{E}[(-sf_0) \div f|E] \implies \|\Psi_f^*(\delta)^-\| \leq 2\varepsilon \|f_0\|.$$

A similar argument works for $t \nearrow 1$. ■_{3.10}

Remark 3.11. If one is bothered with the presence of f_0 in the definition, one can also prove that

$$\Psi_f^*(t) = \sup_{g \in E} tg - \mathbf{E}[g \dot{-} f|E], \quad \mathbf{E}[g \dot{-} f|E] = \sup_{t \in [0,1]} tg - \Psi_f^*(t),$$

which looks formally closer to the (double) scalar Legendre transform. By the reasoning of Lemma 3.10, one can also prove a version of (1), namely

$$\Psi^*(t) = tg - \mathbf{E}[g \dot{-} f] \iff D_t^- \Psi_f^* \leq g \leq D_t^+ \Psi_f^*.$$

This was indeed the path followed in an earlier version of this article, but the present seems more direct.

For $0 < t < 1$, let us define $f_t = D_t^- \Psi^*$, which determines an increasing sequence $(f_t) \subseteq E$. Following Convention 3.8, we may then define (measurably) $\hat{f}: (\Omega \cup \{+, -\}) \times [0, 1] \rightarrow \mathbf{R}$ by $\hat{f}(\omega, t) = f_t(\omega)$, $\hat{f}(+, \cdot) = \|f^+ \upharpoonright_{E^\perp}\|$, $\hat{f}(-, \cdot) = -\|f^- \upharpoonright_{E^\perp}\|$ (ignoring $t = 0, 1$).

Lemma 3.12. *With $E \subseteq E'$, f_t and \hat{f} as above, we have $\hat{f} \in E'$, $\text{tp}(f/E) = \text{tp}(\hat{f}/E)$, and $\|f_t\| \leq \frac{\|f\|}{(t-t^2)^{1/p}}$ for all $0 < t < 1$.*

Proof. Let us show that the conditional distribution of $\hat{f} \upharpoonright_E = \hat{f} \cdot \mathbf{1}_\Omega$ with respect to the Borel algebra of Ω , call it Σ , is the same as that of $f \upharpoonright_E$. Indeed, let $f_0 = \mathbf{1}_A$. Then $\hat{f} \mathbf{1}_{\Omega \setminus A} = 0 = f \mathbf{1}_{\Omega \setminus A}$, while for (almost all) $\omega \in A$ we have by Lemma 3.9 and (1):

$$\mathbf{P}[\hat{f} \leq x | \Omega](\omega) = \sup\{t: D_t^- \Psi_f^*(\omega) \leq x\} = \sup\{t: t \leq D_x^+ \Psi_f(\omega)\} = \mathbf{P}[f \leq x | \Omega](\omega).$$

It follows that $\|\hat{f} \upharpoonright_E\|_p = \|f \upharpoonright_E\|_p < \infty$, so $\hat{f} \in E'$, and as in Lemma 3.9 we conclude that the types are the same. It follows that

$$\|f\|_p^p \geq \|\hat{f} \upharpoonright_E\|_p^p = \int_0^1 \|f_t\|_p^p dt = \int_0^1 \|f_t^+\|_p^p + \|f_t^-\|_p^p dt.$$

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

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

whence the stated inequality. ■_{3.12}

A comparison with [BBH11, Section 5] yields that f_t agrees with the conditional slice $\mathbf{S}_{1-t}(f/E)$, as defined there, for all t at which Ψ_f^* is differentiable (namely, at which $D_t^+ \Psi_f^* = D_t^- \Psi_f^* = f_t$), and in particular for almost all t . The realisation constructed above is called there an ‘‘increasing realisation’’, and is shown to be unique when $f = f \upharpoonright_E$. It follows that each 1-type over E has a unique realisation in $f \in E'$ such that $f(\omega, t)$ is increasing in t on $\Omega \times [0, 1]$, is a non negative constant on $\{+\} \times [0, 1]$, and is a non positive constant on $\{-\} \times [0, 1]$. In other words, the map Ψ_f^* (together with $\|f^\pm\|$) not only determines $\text{tp}(f/E)$, but also provides the means for a fairly direct construction of a canonical realisation of the type.

The continuity of Ψ_f^* implies that $\Psi_f^*(t) = \int_0^t f_s ds$. Since $\Psi_f^*(1) = \int_0^1 f_s ds = \mathbf{E}[f|E]$, we may think of $\Psi_f^*(t)$ as a ‘‘part’’ of the conditional expectation, (with complement $\mathbf{E}[f|E] - \Psi_f^*(t) = \int_t^1 f_s ds$), calling it a *partial conditional expectation*, denoted accordingly $\mathbf{E}_t[f|E] = \Psi_f^*(t)$. In particular, $\mathbf{E}_0[f|E] = 0$ and $\mathbf{E}_1[f|E] = \mathbf{E}[f|E]$.

Lemma 3.13. *Assume that $p > 1$. Then $\text{Cb}(f/E) \mapsto \mathbf{E}_t[f|E]$ is definable in $ALpL$ for all $t \in [0, 1]$.*

Proof. For $t = 0, 1$ this is already known by Lemma 3.6. Since $\text{Cb}(f/E) \mapsto \text{Cb}(|f|/E)$ is also definable, $f_0 = \mathbf{E}[|f|/E]$ is uniformly definable from the canonical base. We may therefore assume that $0 < t < 1$, and use f_0 as a parameter. Similarly, for each x the map $\text{Cb}(f/E) \mapsto \Psi_f(x)$ is definable.

First of all, let us fix $N > 0$, and let

$$g_{N,n} = \bigvee_{|k| \leq n} tNk f_0/n - \Psi_f(Nk/n), \quad h_N = \sup_{-N \leq x \leq N} tx f_0 - \Psi_f(x).$$

Then the map $\text{Cb}(f/E) \mapsto g_{N,n}$ is definable (for each n), and $0 \leq h_N - g_{N,n} \leq 2Nf_0/n$, so $g_{N,n} \rightarrow h_N$ at a rate which depends on $\|f\|$ and on N alone. It follows that $\text{Cb}(f/E) \mapsto h_N$ is definable.

By (1) we have $h_N(\omega) = \mathbf{E}_t[f|E](\omega)$ when $|f_t(\omega)| \leq N$, and otherwise $|h_N(\omega) - \mathbf{E}_t[f|E](\omega)| \leq 2$, since

$$-f_0 \leq -\mathbf{E}[f^-|E] \leq h_N \leq \mathbf{E}_t[f|E] \leq t\mathbf{E}[f|E] \leq f_0.$$

The measure of $\{\omega: |f_t(\omega)| > N\}$ vanishes as $N \rightarrow \infty$, and at a rate which depends only on $\|f_t\|$, which we already bounded in terms of t and $\|f\|$. We conclude that $h_N \rightarrow \mathbf{E}_t[f|E]$ at a rate which depends only on $\|f\|$ and t , which is enough. $\blacksquare_{3.13}$

Theorem 3.14. *For every $p \in (1, \infty)$ and every dense subset $D \subseteq (0, 1)$ (e.g., $D = \mathbf{Q} \cap (0, 1)$), the tuple $(\|f^+\|, \|f^-\|, \mathbf{E}_t[f|E])_{t \in D}$ is a uniform canonical base for $\text{tp}(f/E)$, and $(\mathbf{E}_t[f|E])_{t \in D}$ is a weakly uniform canonical base in the home sort.*

Proof. By Proposition 1.12 it is enough to prove the first assertion. We have already seen that $\text{Cb}(f/E) \mapsto \mathbf{E}_t[f|E]$ is definable, and clearly $\text{Cb}(f/E) \mapsto \|f^\pm\|$ are. On the other hand, by Lemma 3.10 the tuple $(\|f^+\|, \|f^-\|, \mathbf{E}_t[f|E])_{t \in D}$ determines Ψ_f and therefore, by Lemma 3.9, it determines $\text{tp}(f/E)$. $\blacksquare_{3.14}$

We have already observed that $\mathbf{E}_t[f|E] \rightarrow 0$ as $t \rightarrow 0$ and $\mathbf{E}_t[f|E] \rightarrow \mathbf{E}[f|E]$ as $t \rightarrow 1$. Moreover, for $p > 1$ (fixed) the rate of convergence depends uniformly on $\|f\|$. Indeed, otherwise Lemma 3.13 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 $\mu(\Omega) = 1$ and $f_\varepsilon(x, t) = -\varepsilon^{-1} \mathbf{1}_{\Omega \times [0, \varepsilon]}$. Then $\|f_\varepsilon\| = 1$ and $\mathbf{E}_\varepsilon[f_\varepsilon|E] = -\mathbf{1}_\Omega$, also of norm one, so the rate of convergence is not uniform. Thus Lemma 3.13, and therefore Lemma 3.6, 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 3.14 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 3.15. *Let $p \in [1, \infty)$. Then for every $0 < s < t < 1$, the map $\text{Cb}(f/E) \mapsto \mathbf{E}_{[t,s]}[f|E]$ 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 $\text{Cb}(f/E)$ (as in the proof of Lemma 3.3, by $\mathbf{E}_{[t,s]}[f|E]^{L_q}$ we mean $\mathbf{E}_{[t,s]}[f|E]$ as calculated in the L_q lattice defined in the L_1 lattice). It will be enough to show that as $q \rightarrow 1$, $\mathbf{E}_{[t,s]}[f^{1/q}|E]^q \rightarrow \mathbf{E}_{[t,s]}[f|E]$ at a rate which only depends on $\|f\|$.

Concentrating on what happens over a single $\omega \in \Omega$, all functions in E become constants, and we may identify f with the function $f(t) = f_t$, which is increasing on $[0, 1]$. If $f_0 = 0$ then everything is zero, so we may assume that $f_0 = \int_0^1 |f(t)| dt = 1$. Since f is increasing, we must have $f(r) \in [-\frac{1}{t}, \frac{1}{1-s}]$ for all $r \in [t, s]$, and by, say, dominated convergence,

$$\left(\int_t^s f(r)^{\frac{1}{q}} dr \right)^q \rightarrow_{q \rightarrow 1} \int_t^s f(r) dr.$$

The space of all increasing functions $[t, s] \rightarrow [-\frac{1}{t}, \frac{1}{1-s}]$ is compact in the topology of convergence in measure, which means that the convergence above is uniform in f , that is to say that for any desired $\varepsilon > 0$ there exists $q_0 = q_0(t, s, \varepsilon) > 1$ 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 we obtain

$$|\mathbf{E}_{[t,s]}[f^{1/q}|E]^q - \mathbf{E}_{[t,s]}[f|E]| < \varepsilon \|f\|,$$

as desired. $\blacksquare_{3.15}$

Theorem 3.16. *For every $p \in [1, \infty)$ and every dense subset $D \subseteq (0, 1)$ (e.g., $D = \mathbf{Q} \cap (0, 1)$), the tuple $(\|f^+\|, \|f^-\|, \mathbf{E}_{[t,s]}[f|E])_{t,s \in D, t < s}$ is a uniform canonical base for $\text{tp}(f/E)$, and $(\mathbf{E}_{[t,s]}[f|E])_{t,s \in D, t < s}$ is a weakly uniform canonical base in the home sort.*

Proof. As for Theorem 3.14, since $\mathbf{E}_s[f|E] = \lim_{t \rightarrow 0} \mathbf{E}_{[t,s]}[f|E]$. $\blacksquare_{3.16}$

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.

Fact 3.17. *The n -type $\text{tp}(\bar{f}|_E/E)$ is determined by the 1-types $\text{tp}(\bar{k} \cdot \bar{f}/E)$, 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 $\text{tp}(\bar{t} \cdot \bar{f}/E)$, and in particular $\text{tp}(\bar{t} \cdot \bar{f}|_E/E)$ for all $\bar{t} \in \mathbf{Q}^n$ and therefore for all $\bar{t} \in \mathbf{R}^n$. Now apply [BBH11, Proposition 3.7]. ■_{3.17}

Theorem 3.18. *Let Cb be a uniform canonical base map for 1-types. Then*

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

is a uniform canonical base map for n -types. (We may view $\text{tp}(\bar{f})$ as a sequence in $[0, 1]$ via any embedding of $\mathbf{S}_n(\text{ALpL})$ in $[0, 1]^{\aleph_0}$.)

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

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

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

Proof. For the first assertion, it is enough to show that $\text{tp}(\bar{f}/E)$ is determined by these data. Indeed, $\text{tp}(\bar{f}|_E/E)$ is already known to be determined. Let $h = \sum_i |f_i|$. Then $\text{tp}(\bar{f}|_E/E)$ determines $\text{tp}(h|_E/E)$, and in particular $\|h|_E\|$, while $\|h\|$ is determined by $\text{tp}(\bar{f})$, so $\|h|_{E^\perp}\|$ is determined as well. Alongside the facts that $h|_{E^\perp}$ is positive and orthogonal to E , this is enough to determine $\text{tp}(h|_{E^\perp}/E)$, so $\text{tp}(h/E)$ is determined. We may choose any realisation of this type, in a sufficiently saturated extension of E , and without changing $\text{tp}(\bar{f}/E)$, we may assume this realisation is indeed h .

We may further assume that $h|_E = \mathbf{1}_A$ and $h|_{E^\perp} = \mathbf{1}_B$ in some concrete presentation of the ambient model. Now $\text{tp}(\bar{f})$ determines $\text{tp}(\bar{f}/h)$, which, again by [BBH11, Proposition 3.7], can be identified with the joint conditional distribution of \bar{f} with respect to $\{\emptyset, A \cup B\}$ (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, $\text{tp}(\bar{f}|_E/E)$ determines $\text{tp}(\bar{f}|_E)$ and thus $\text{tp}(\bar{f}|_E/h|_E)$, which can be identified with the joint conditional distribution of $\bar{f}|_E$ with respect to $\{\emptyset, A\}$. Subtracting, we obtain the joint conditional distribution of $\bar{f}|_{E^\perp}$ with respect to $\{\emptyset, B\}$, namely $\text{tp}(\bar{f}|_{E^\perp}/h|_{E^\perp})$, and thus $\text{tp}(\bar{f}|_{E^\perp})$ and finally $\text{tp}(\bar{f}|_{E^\perp}/E, \bar{f}|_E)$. Thus $\text{tp}(\bar{f}/E)$ is known and the proof of the first assertion is complete. The second assertion follows. ■_{3.18}

Corollary 3.19. *For every $p \in [1, \infty)$ and every dense subset $D \subseteq (0, 1)$ the tuples*

$$(\text{tp}(\bar{f}), \mathbf{E}_{[t,s]}[\bar{k} \cdot \bar{f}|E])_{t,s \in D, t < s, \bar{k} \in \mathbf{Z}^n}$$

is a uniform canonical base for $\text{tp}(\bar{f}/E)$, and

$$(\mathbf{E}_{[t,s]}[\bar{k} \cdot \bar{f}|E])_{t,s \in D, t < s, \bar{k} \in \mathbf{Z}^n}$$

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 3.20. 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 $\text{tp}(\bar{f})$. Indeed, let f_1, f_2 and f_3 be disjoint positive functions of norm one, and let

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

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

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 $(\text{Cb}_n)_n$, 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}_{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 $f_{\text{Cb}}(\bar{x})$, where $n = |\bar{x}|$. In the continuous setting, uniform continuity moduli for the f_{Cb} are as per Lemma 1.7. We let $\mathcal{L}_{P, \text{Cb}} = \mathcal{L}_P \cup \mathcal{L}_{\text{Cb}}$.

By a *pair* of models of T we mean any elementary extension $\mathcal{N} \preceq \mathcal{M} \models T$. We shall identify such a pair with the structures (\mathcal{M}, P) , $(\mathcal{M}, f_{\text{Cb}})$ or $(\mathcal{M}, P, f_{\text{Cb}})$, as will be convenient, where $P(x) = d(x, N)$ and $f_{\text{Cb}}(\bar{x}) = \text{Cb}(\bar{x}/N)$. 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, and we shall denote its theory by $T_{P,0}$. Similarly, $T_{P,\text{Cb},0}$ will be the $\mathcal{L}_{P,\text{Cb}}$ -theory of pairs, which consists in addition of the axioms saying that $f_{\text{Cb}}(\bar{x}) = \text{Cb}(\bar{x}/P)$.

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

$$P(x) = \exists y d[x = y](y, \text{Cb}(x/P)), \quad P(x) = \inf_y d[d(x, y)](y, \text{Cb}(x/P)).$$

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}_{Cb} -theory $T_{\text{Cb},0}$, for which $T_{P,\text{Cb},0}$ 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 $T_{\text{Cb},0}$ admits amalgamation over arbitrary sets (which generate isomorphic substructures). If T is complete, then $T_{\text{Cb},0}$ also admits the joint embedding property.*

Proof. Let $(\mathcal{M}_i, P) \models T_{\text{Cb},0}$ for $i = 0, 1$ and let A be a common sub-structure. Since, for every sentence φ , we have $\varphi^{\mathcal{M}_i} = d\varphi(\emptyset, \text{Cb}(\emptyset/P))$ and $\text{Cb}(\emptyset/P) \in A$, we have $\mathcal{M}_0 \equiv \mathcal{M}_1$ even if T is incomplete. Let \mathcal{N} be a sufficiently saturated model of the common complete theory.

We first choose an embedding $\theta: A \rightarrow N$. Then, we may choose two extensions, $\theta_i: P(\mathcal{M}_i) \cup A \rightarrow N$, such that in addition $\theta_0 P(\mathcal{M}_0) \downarrow_{\theta A} \theta_1 P(\mathcal{M}_1)$. Since A is a sub-structure in \mathcal{L}_{Cb} , we know that $A \downarrow_{P(A)} P(\mathcal{M}_i)$, and the same holds for the images under θ_i . It follows that $\theta A \downarrow_{\theta P(A)} \theta_0 P(\mathcal{M}_0), \theta_1 P(\mathcal{M}_1)$, and we may choose a sub-model $P(\mathcal{N}) \preceq \mathcal{N}$ such that $P(\mathcal{N}) \supseteq \theta_i P(\mathcal{M}_i)$ and $\theta A \downarrow_{\theta P(A)} P(\mathcal{N})$. We may now extend each θ_i to all of \mathcal{M}_i , such that $\theta_i \mathcal{M}_i \downarrow_{\theta A, \theta_i P(\mathcal{M}_i)} P(\mathcal{N})$. Then in particular $\theta_i \mathcal{M}_i \downarrow_{\theta_i P(\mathcal{M}_i)} P(\mathcal{N})$. Then $\theta_i: (\mathcal{M}_i, P) \rightarrow (\mathcal{N}, P)$ respects \mathcal{L}_{Cb} , and we are done.

If T is complete then we can amalgamate models over the (unique) \emptyset -generated substructure. ■_{4.1}

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

It follows that a model companion of $T_{\text{Cb},0}$ (or of $T_{P,\text{Cb},0}$, 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 $T_{\text{Cb},0}$, the restriction of every quantifier-free \mathcal{L}_{Cb} -formula to P is \mathcal{L} -definable there.*

Proof. This follows immediately from the facts that, first, for $\bar{a} \in P$ we have $\text{Cb}(\bar{a}/P) = \text{Cb}(\bar{a}/\bar{a})$, and second, the map $\bar{a} \mapsto \text{Cb}(\bar{a}/\bar{a})$ is definable in \mathcal{L} . ■_{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} (namely, the $\mathcal{L}(\mathcal{M})$ -structure $(\mathcal{N}, a)_{a \in \mathcal{M}}$ is approximately \aleph_0 -saturated). We define T_{Cb}^b to be the \mathcal{L}_{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 img Cb_n is definable for each n (i.e., its projection to every finite sort is a definable set in that sort).*
- (ii) *If $(\mathcal{M}, P) \models T_{\text{Cb}}^b$ is κ -saturated for some $\kappa > |T|$ then \mathcal{M} is κ -saturated over $P(\mathcal{M})$. In particular, every sufficiently saturated model of T_{Cb}^b is itself a beautiful pair.*
- (iii) *The theory $T_{\text{Cb},0}$ admits a companion T_{Cb} such that for some $\kappa > |T|$, if $(\mathcal{M}, P) \models T_{\text{Cb}}$ is κ -saturated then \mathcal{M} is κ -saturated over $P(\mathcal{M})$. Moreover, such a companion is necessarily the model companion of $T_{\text{Cb},0}$.*
- (iv) *The theory $T_{\text{Cb},0}$ 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 $(\mathcal{M}, P) \models T_{\text{Cb},0}$, 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) \implies (ii). Let T_{Cb} consist of $T_{\text{Cb},0}$ along with the axioms saying that for every canonical base $Z \in \text{img Cb}$ 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 img Cb is definable, this axiom can indeed be expressed, and clearly $T_{\text{Cb}} \subseteq T_{\text{Cb}}^b$.

Now let $(\mathcal{M}, P) \models T_{\text{Cb}}$ 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} , $C = \text{Cb}(\bar{a}/M) \in M$, and let $A \subseteq M$, $|A| < \kappa$. Let $\pi(\bar{x})$ be the partial \mathcal{L}_{Cb} -type saying that \bar{x} realises $\text{tp}(\bar{a}/A \cup P)$, i.e., that

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

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) \implies (iii). It is easy to check that $T_{\text{Cb},0}$ and T_{Cb}^b 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 $p(\bar{x}, \bar{y}) = \text{tp}(\bar{a}, \bar{b}/P(N))$, $q(\bar{y}) = \text{tp}(\bar{b}/P(N))$, $C = \text{Cb}(p)$, $D = \text{Cb}(q)$, observing that then $D \subseteq P(M)$. 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}_{Cb} -type over \bar{b} , which concludes the proof.

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

(iv) \implies (i). For $(\mathcal{M}, P) \models T_{\text{Cb}}$, let $(\text{img Cb}_n)^P$ denote the set of all $\text{Cb}(p)$ as $p \in S_n(P(M))$ (namely, the type-definable set img Cb_n as interpreted in the structure $P(M)$). In a sufficiently saturated model of T_{Cb} , $(\text{img Cb}_n)^P$ is exactly $f_{\text{Cb}}(M^n)$, which is definable as a definable image of a definable set. It follows that it is uniformly definable in models of T_{Cb} (and the image of f_{Cb} is always dense there). By quantifier elimination, it is quantifier-free definable. It follows that $(\text{img Cb}_n)^P$ is \mathcal{L} -definable in P , which means exactly that img Cb_n is definable.

(iv) \implies (v). A close inspection of the argument for (ii) \implies (iii) reveals that it proves the following intermediary result: if $(\mathcal{M}, P) \models T_{\text{Cb},0}$, $P(\mathcal{M})$ 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) \implies (vi). Clear.

(vi) \implies (i). Here we also assume that T is complete. For a formula $\varphi(\bar{x}, \bar{y})$, with $|\bar{x}| = n$, let f_φ be the definable function sending Z in the target sort of Cb_n to the canonical parameter of $d\varphi(\bar{y}, Z)$. The target sort of f_φ is equipped with a canonical definable distance $d(z, w) = \sup_{\bar{y}} |d\varphi(\bar{y}, z) - d\varphi(\bar{y}, w)|$. We may define $\text{Cb}_\varphi = f_\varphi \circ \text{Cb}_n$, and $\text{Cb}_\varphi(p)$ is a (uniform) φ -canonical base for p . Through the coding of several formulae in one, we see that img Cb_n is definable if and only if img Cb_φ is for every such φ . We may therefore assume that img Cb_φ is not definable for some φ , keeping in mind that it is always type-definable – we may in fact identify it with a closed set in $S_z(T)$.

Non definability then implies that for some $\varepsilon > 0$, the set of types $q(z)$ such that $d(q, \text{img Cb}_\varphi) \geq \varepsilon$ (namely, such that $q(z) \cup r(z) \vdash d(z, w) \geq \varepsilon$ for every $r \in \text{img Cb}_\varphi$) has an accumulation point $r \in \text{img Cb}_\varphi \subseteq S_z(T)$. The partial type asserting that $z \in P$, $r(z)$, and $\inf_{\bar{x}} g(z, f_\varphi \circ f_{\text{Cb}}(\bar{x})) \geq \varepsilon$ is then approximately finitely realised in (\mathcal{M}, P) , and therefore, by the saturation assumption, realised there, say by c . Then $d\varphi(\bar{y}, c)$ defines a complete φ -type p over P , which is realised in \mathcal{M} , by the saturation assumption for \mathcal{M} over P , so $d(c, f_\varphi \circ f_{\text{Cb}}(M^n)) = 0$, a contradiction. ■_{4.4}

Corollary 4.5. *The image sets img Cb_n are definable if and only if $T_{\text{Cb},0}$ admits a model companion, if and only if this model companion is T_{Cb}^b (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_{Cb}^b is the model completion of $T_{\text{Cb},0}$.*

Proof. In an \aleph_0 -categorical theory every type-definable set is definable. ■_{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 $T = \text{ACMVF}$.*

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 $Q(X, \bar{e}, f)$ where $f \in P$ (or more exactly, of the homogenisation thereof, so ∞ can also be obtained a root). Let us define T_{Cb} to consist of $T_{\text{Cb},0}$ along with the axioms saying that every $r \in [0, 1]$, $b \in P$ and $Q(X, \bar{e}, \bar{V})$ as above there is a such that $d(a, b) = d(a, \sqrt{Q(X, \bar{e}, P)}) = r$, 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 $\{x^{-1} : x \in \sqrt{Q(X, \bar{e}, P)}\} = \sqrt{Q'(X, \bar{e}, P)}$, we may assume that $|b| \leq 1$. We may then add a such that $|a - c| = \max r, |b - c|$ for all $c \in \sqrt{Q(X, \bar{e}, P)} \setminus \{\infty\}$, so in particular $|a| = \max\{|b|, r\} \leq 1$. It follows that $d(a, \sqrt{Q(X, \bar{e}, P)}) = d(a, b) = r$.

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 $b \in P \setminus \{\infty\}$, 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 of $(K\mathbf{P}^1, P)$) there exists x such that $d(a, B) = d(a, b) = r$, and as in the previous paragraph this x is as desired. Otherwise, $|b| > 1$. Again by saturation, we find a such that for all $c \in B \setminus \{\infty\}$, $|a - c| = \max r', |z - b^{-1}|$ 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 $c \in B \setminus \{\infty\}$, $|a^{-1} - c| = \max r \left| \frac{c}{b} \right|, |b - c|$ (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 $p(x) \in S_1(PA)$, and let q be any global extension of p to $K\mathbf{P}^1$. Since $K\mathbf{P}^1$ is \aleph_1 -saturated, there exists $b \in K\mathbf{P}^1$ such that $r = d(x, b)^{q(x)} = d[d(x, y)](b, \text{Cb}(q))$ is minimal. Replacing p with $q|_{PA}$, we may assume that $b \in A$. By our previous claim, there exists $a \in K\mathbf{P}^1$ 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 $(K\mathbf{P}^1, P)$ of the companion T_{Cb} , $K\mathbf{P}^1$ is \aleph_1 -saturated over P , so by Theorem 4.4 T_{Cb} is the model completion of $T_{\text{Cb},0}$, 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 $\text{dcl}(S) = \text{dcl}(\emptyset)$. 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 $[a_r] = [b_s]$.) 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

- [BBH] Itai BEN YAACOV, Alexander BERENSTEIN, and C. Ward HENSON, *Almost indiscernible sequences and convergence of canonical bases*, submitted, arXiv:0907.4508.
- [BBH11] ———, *Model-theoretic independence in the Banach lattices $L_p(\mu)$* , Israel Journal of Mathematics **183** (2011), 285–320, doi:10.1007/s11856-011-0050-4, arXiv:0907.5273.
- [Bena] Itai BEN YAACOV, *Model theoretic properties of metric valued fields*, submitted, arXiv:0907.4560.
- [Benb] ———, *On theories of random variables*, Israel Journal of Mathematics, to appear, arXiv:0901.1584.
- [Ben06] ———, *Schrödinger's cat*, Israel Journal of Mathematics **153** (2006), 157–191, doi:10.1007/BF02771782.
- [Ben09] ———, *Modular functionals and perturbations of Nakano spaces*, Journal of Logic and Analysis **1:1** (2009), 1–42, doi:10.4115/jla.2009.1.1, arXiv:0802.4285.
- [BPV03] Itai BEN YAACOV, Anand PILLAY, and Evgueni VASSILIEV, *Lovely pairs of models*, Annals of Pure and Applied Logic **122** (2003), 235–261, doi:10.1016/S0168-0072(03)00018-6.
- [BU10] Itai BEN YAACOV and Alexander USVYATSOV, *Continuous first order logic and local stability*, Transactions of the American Mathematical Society **362** (2010), no. 10, 5213–5259, doi:10.1090/S0002-9947-10-04837-3, arXiv:0801.4303.
- [Fre04] D. H. FREMLIN, *Measure theory volume 3: Measure algebras*, Torres Fremlin, 25 Ireton Road, Colchester CO3 3AT, England, 2004.
- [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), no. 1, 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.
- [Roc70] R. Tyrrell ROCKAFELLAR, *Convex analysis*, Princeton Mathematical Series, No. 28, Princeton University Press, Princeton, N.J., 1970.
- [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/>