On the model theory of group actions on probability measure algebras

C. Ward Henson
University of Illinois
(collaboration with Alex Berenstein)

Octrober 2, 2018
Notre Dame University

The continuous signatures we use

- L^{pr}: signature for probability measure algebras:
  - Boolean operations: $\cup, \cap, \cdot^c$
  - Constants: $0, 1$
  - Measure: μ (a unary predicate; values in $[0,1]$)
  - Metric: d (a binary predicate, values in $[0,1]$)
  - specify suitable moduli of uniform continuity; for example $|\mu(x) - \mu(y)| \leq d(x,y)$.

- L^F: signature for measure algebras with automorphisms
  - F is a set of unary function symbols to be added to L^{pr}
  - For each $f \in F$, require $d(f(x), f(y)) \leq d(x,y)$.

probability measure algebras

Let $(X, \mathcal{B}, \mu)$ be a probability space; *i.e.*, X is a nonempty set, $\mathcal{B}$ is a σ-complete boolean algebra of subsets of X that contains $\varnothing$ and X, and $\mu\colon \mathcal{B} \to [0,1]$ is a σ-additive probability measure on $\mathcal{B}$.

The measure algebra of $(X, \mathcal{B}, \mu)$ is the L^{pr}-structure $(\mathcal{B}/\mathcal{I}, \mu, d)$.

$$\text{Here } \mathcal{I} = \{a \in \mathcal{B} \mid \mu(a) = 0\} \text{ is the ideal of null sets,}$$

$$\mu([a]_{\mathcal{I}}) := \mu(a) \text{ for all } a \in \mathcal{B}, \text{ and}$$

$$d([a]_{\mathcal{I}}, [b]_{\mathcal{I}}) := \mu(a \triangle b) \text{ for all } a, b \in \mathcal{B}.$$

We axiomatize the class of probability measure algebras using continuous logic.

Pr: L^{pr}-axioms for probability measure algebras

- (Boolean algebra axioms)

$$\sup_{x_1} \ldots \sup_{x_n} \; d(t_1(x), t_2(x)) \; = \; 0$$

  for various boolean terms t_1, t_2.

- (Measure axioms)

$$\mu(0) = 0 \text{ and } \mu(1) = 1 \text{ and}$$
$$\sup_x \sup_y \big|[\mu(x \cup y) + \mu(x \cap y)] - [\mu(x) + \mu(y)]\big| = 0.$$

- (Connection between d and μ)

$$\sup_x \sup_y |d(x, y) - \mu(x \Delta y)| = 0.$$

  Here $x \Delta y$ denotes the symmetric difference:

$$x \Delta y = (x \cap y^c) \cup (y \cap x^c).$$

APr: atomless probability measure algebras

- $APr = Pr$ together with:

$$\sup_x \inf_y \left|\mu(x \cap y) - \mu(x \cap y^c)\right| = 0.$$

- Note that Pr is an $\forall$ theory and APr is an $\forall\exists$ theory.

Theorem

- *The models of* Pr *are exactly the measure algebras of probability spaces.*
- *The models of* APr *are exactly the measure algebras of atomless probability spaces.*
- *Therefore* $(APr)_\forall = Pr$.

Theorem

- *The theory APr admits quantifier elimination and is complete.*
- *Hence APr is the model companion of Pr.*
- *Every model of APr is ω-saturated.*
- *Equivalently, if $M \vDash APr$ and $a \in M^n$, then $\mathrm{Th}(M, a)$ is separably categorical.*

Corollary

For every L^{pr}-formula $\theta(x_1, \ldots, x_n)$, the zeroset of θ^M is a definable set $\subseteq M^n$ (uniformly) in all models M of APr.

some notation

Let $M = (\mathcal{B}, \mu, d) \models Pr$ and $C \subseteq M$.

- Let $C^{\#}$ denote the subalgebra of $\mathcal{B}$ generated by C.
- Let $\langle C \rangle$ denote the σ-subalgebra of $\mathcal{B}$ generated by C.
- Note that $\langle C \rangle$ is the closure of $C^{\#}$ in the metric space $(\mathcal{B}, d)$.
- Also, if C is finite, then $C^{\#}$ is also finite and $\langle C \rangle = C^{\#}$.

more notation

Again let $M = (\mathcal{B}, \mu, d) \models Pr$.

- For $a \in M^n$, let a^* denote the 2^n-tuple of elements of the form $a_1^{e_1} \cap \cdots \cap a_n^{e_n}$, where $e_1, \ldots, e_n$ is any sequence from $\{-1, +1\}$, with the cordinates of a^* written in lexicographic order of $(e_1, \ldots, e_n)$. (Here $a_i^{+1} := a_i$ and $a_i^{-1} := a_i^c$.) Note that $a*$ is a partition of 1 in $\mathcal{B}$, and that each a_i is the union of those coordinates of a^* with $e_i = 1$.
- We define x^* in the same way when $x = x_1, \ldots, x_n$ is a tuple of distinct variables.

Description of types in *APr*

Proposition

Let $M = (\mathcal{B}, \mu, d) \models APr$*. Consider* $a, b \in M^m$ *and* $c \in M^n$*. The following are equivalent:*

(1) $\mathrm{tp}(a/c) = \mathrm{tp}(b/c)$*;*

(2) $\mathrm{tp}(a^*/c^*) = \mathrm{tp}(b^*/c^*)$*;*

(3) For each $i = 1, \ldots, 2^m$ *and* $j = 1, \ldots, 2^n$*, we have*

$$\mu(a_i^* \cap c_j^*) = \mu(b_i^* \cap c_j^*);$$

(4) For each $i = 1 \ldots, 2^m$ *, we have* $\mathrm{tp}(a_i^*/c) = \mathrm{tp}(b_i^*/c)$*;*

From (3) we get an L^{pr}*-formula* $\varphi(x, y, z)$ *such that*

$$\{(a, b, c) \in M^{2m+n} \mid tp(a/c) = tp(b/c)\}$$

is the zeroset of φ^M *in all models* M *of* APr*.*

Hence this is a definable set, uniformly in M*.*

Distance between types of partitions of 1 in *APr*

Proposition

Let $M = (\mathcal{B}, \mu, d) \models APr$*. Consider* $a, b \in M^m$ *and* $c \in M^n$*, with each of* a, b, c *being a partition of* 1*. Then*

$$d\big(\operatorname{tp}(a/c), \operatorname{tp}(b/c)\big) = \max_i \sum_j \big|\mu(a_i \cap c_j) - \mu(b_i \cap c_j)\big|.$$

Corollary

Let $M = (\mathcal{B}, \mu, d) \models APr$*. Consider* $a, b \in M^m$ *and* $C \subseteq M$*, with each of* a, b *being a partition of* 1*. Then*

$$d\big(\operatorname{tp}(a/C), \operatorname{tp}(b/C)\big) = \max_i \| \mathbb{P}(a_i \mid \langle C \rangle) - \mathbb{P}(b_i \mid \langle C \rangle) \|_1.$$

Stability and non-dividing in *APr*

It follows from the preceding result that APr is ω-stable; that is, for each countable set $C \subseteq M \models APr$, the type spaces $S_n(C)$ are separable. Hence there is a well-behaved concept of non-dividing for types in APr, which can be described as follows:

Proposition

Let $D \subseteq C \subseteq M = (\mathcal{B}, \mu, d) \models APr$ and $a \in M^n$. Then the following are equivalent:
(1) $\mathrm{tp}(a/C)$ *does not divide over* D.
(2) $tp(a^*/C)$ does not divide over D.
(3) For every $j = 1, \ldots, 2^n$, we have
$$\mathbb{P}(a_j^* \mid \langle C \rangle) = \mathbb{P}(a_j^* \mid \langle D \rangle).$$

It follows that every type p in $S_n(D)$ is stationary; i.e., p has a unique non-dividing extension to a type over C.

A calculation based on preceding results yields:

Corollary

Suppose $M = (\mathcal{B}, \mu, d) \models APr$, and that $u \in M^r, v \in M^s$, and $w \in M^t$ are each partitions of 1. Suppose $q \in S_r(vw)$ is the unique non-dividing extension of $\mathrm{tp}(u/w)$. Then

$$d(\mathrm{tp}(u/vw), q) = \max_i \sum_{j,k} \left| \mu(u_i \cap v_j \cap w_k) - \frac{\mu(u_i \cap w_k)\mu(v_j \cap w_k)}{\mu(w_k)} \right|.$$

This gives an L^{pr}-formula $\psi(x, y, z)$ such that

$\{(a, c, d) \in M^{m+n+p} \mid tp(a/cd) \text{ is non-dividing over } d\} =$
$\{(a, c, d) \in M^{m+n+p} \mid tp(a^*/c^*d^*) \text{ is non-dividing over } d^*\}$

is the zeroset of ψ^M in all models M of APr.

Hence this is a definable set, uniformly in M.

Main topic: actions of groups

Given a set F, recall that L^F is the signature obtained by adding the elements of F to L^{pr} as unary function symbols.

Let T^F denote the L^F-theory obtained by adding to Pr conditions expressing that each $f \in F$ is interpreted by an automorphism of the underlying model of Pr.

The needed conditions express that f preserves the boolean operations and the measure (hence also the metric) and that f is surjective. This last requirement is expressed by

$$\sup_y \inf_x d(y, f(x)) = 0.$$

For a word $w = f_1 \cdots f_n$ on the alphabet F and a variable x, we set

$$w(x) := f_1(\cdots(f_n(x))\ldots).$$

Further, if $x = x_1, \ldots, x_n$ is a sequence of variables, we set

$$w(x) := (w(x_1), \ldots, w(x_n)).$$

Suppose G is a finite or countable discrete group. Let $T(G)$ denote the L^G-theory obtained by adding to T^G the conditions

$$\sup_x d(v(x), w(x)) = 0$$

for every words v, w on the set G such that v and w represent the same element in the group G.

Obviously, models of $T(G)$ correspond exactly to actions of G by automorphisms on probability measure algebras.

Theorem

- *If G is amenable, then $T(G)$ has a model companion $T^*(G)$.*
- *The theory $T^*(G)$ extends APr and it admits quantifier elimination and is complete.*
- *Further, $T^*(G)$ is stable, with its non-dividing relation obtained by applying the non-dividing relation for APr to the $T^*(G)$-definable closure of the sets involved.*
- *Hence, all types over sets are stationary.*

- The main ingredient of the proof is a generalization of Rokhlin's Lemma (which applies to $\mathbb{Z}$-actions) to actions of amenable groups, due to Ornstein and Weiss. It allows us to give a reasonable description of the models of $T^*(G)$.

- For example, $M \vDash T^*(\mathbb{Z})$ iff 1^M is an aperiodic automorphism of the reduct of M to L^{pr}, which is a model of APr.

Theorem

For each set F, the theory T^F has a model companion, which has quantifier elimination and is complete.

To prove this result we need to axiomatize the class of existentially closed (e.c.) models of T^F. Recall that $M \vDash T^F$ is e.c. iff

$$\big(\inf_y \theta(a, y)\big)^M = \big(\inf_y \theta(a, y)\big)^N$$

whenever $\theta(x, y)$ is a quantifier-free L^F-formula, N is a model of T^F with $M \subseteq N$, and a is a tuple from M.

Since T^F is an inductive theory, we know that every model of T^F is contained in an e.c. model.

To avoid notational complexity and make the key ideas more clear, we assume $F = \{f, g\}$.
As a first step in the proof, one shows that a model M of T^F is e.c. iff M satisfies APr and the following statement for all finite tuples a, b, c, d from M:
Assume

- $\mathrm{tp}(b/f(d)) = \mathrm{tp}(f(a)/f(d))$.
- $\mathrm{tp}(c/g(d)) = \mathrm{tp}(g(a)/g(d))$.
- $\mathrm{tp}(a/df(d)g(d))$ is the non-dividing extension of $\mathrm{tp}(a/d)$.
- $\mathrm{tp}(b/df(d)g(d))$ is the non-dividing extension of $\mathrm{tp}(b/d)$.
- $\mathrm{tp}(c/df(d)g(d))$ is the non-dividing extension of $\mathrm{tp}(c/d)$.

Then there exists a tuple e from M such that

$$\mathrm{tp}(ef(e)g(e)/df(d)g(d)) = \mathrm{tp}(abc/df(d)g(d)).$$

(Here "tp" and "non-dividing" are meant in the sense of APr.)

The proof of the preceding claim uses a stability argument that goes back to work of Kikyo and Pillay.

From our discussion of APr, we know that there are L^{pr}-formulas α and β such that the assumptions are true of a, b, c, d in M iff

$$\alpha^M(a, b, c, d, f(a), g(a), f(d), g(d)) = 0$$

and the conclusion holds of a, b, c, d, e in M iff

$$\beta^M(a, b, c, d, e, f(d), g(d), f(e), g(e)) = 0.$$

(Here α and β depend on the lengths of the tuples a, b, c, d, e.)

AND we know that the zerosets of α and β are definable sets in the theory APr.

However, because we are working in continuous model theory, we cannot yet conclude that the desired implications are expressible by axioms. What happens is that the formulas α and β can be chosen (without changing their zerosets) so that every e.c. model M of T^F satisfies the following stronger requirement (which is expressed by a condition in continuous logic):

$$\sup_{a,b,c,d} \big(\inf_e \beta(a,b,c,d,e,f(d),g(d),f(e),g(e)) \leq \alpha(a,b,c,d,f(a),g(a),f(d),g(d)) \big) = 0.$$

These conditions then axiomatize the e.c. models of T^F.