III. Measures and forking in NIP theories

Anand Pillay

University of Leeds

March 25, 2008

1 Preliminaries

Introduction

We work as before in a very saturated model $\bar{M}$ of a complete first order theory M. S denotes some sort (e.g. elements or n-tuples if T is 1-sorted) or even an ambient $\emptyset$-definable set.

- By a Keisler measure on sort S over A, we mean a finitely additive probability measure μ on A-definable subsets of S (or on formulas over A in sort S), namely
- For each A-definable subset X of S, $\mu(X) \in [0,1]$, $\mu(S) = 1$, $\mu(\emptyset) = 0$, and if X, Y are disjoint then $\mu(X \cup Y) = \mu(X) + \mu(Y)$.
- By a global Keisler measure on sort S we mean a Keisler measure on S over $\bar{M}$.
- A special case of a Keisler measure over A is a complete type (in sort S) over A.
- Any Keisler measure over A extends to a global Keisler measure. (Exercise.)
- As an example, let $\bar{M}$ be a big real closed field, containing therefore R as an elementary substructure. Let I be the (nonstandard) interval $[0,1]$ in $\bar{M}$.
- Lebesgue measure on the real unit interval induces a Keisler measure μ on I over R. Moreover μ has a *unique* extension to a global Keisler measure on I. (Exercise.)

We may suppress mention of the ambient sort S, but x will typically denote a variable of that sort.

Lemma 1.1. *A Keisler measure over A is the "same thing" as a regular Borel probability measure on the compact space $S_x(A)$ of complete types over A in variable x.*

- Explanation.
- Regularity of a Borel probability measure β on a compact space C means that for any Borel subset B of C and $\epsilon > 0$ there are open U and closed D such that $D \subseteq B \subseteq U$ and $\beta(U \setminus D) < \epsilon$.
- Note that any Keisler measure μ over A determines a finitely additive probability measure μ on the clopens of $S_x(A)$, and Keisler shows how μ extends to a Borel probability measure β on $S(A)$ which he also shows to be regular.
- On the other hand if β is a Borel probability meaure on $S(A)$ then the restriction of β to the clopens of $S(A)$ gives a Keisler measure over A.
- If β is also regular then for any closed subset D of $S(A)$, $\beta(D)$ will be the infimum of the $\beta(D')$ for D' clopen containing D, hence β is determined by μ.

Basic results

A basic result, left as an exercise, is:

Lemma 1.2. *Suppose μ is a (global) Keisler measure, $(b_i : i < \omega)$ is an indiscernible sequence, $\phi(x, y) \in L$ and for some $\epsilon > 0$, $\mu(\phi(x, b_i) \geq \epsilon$ for all i. Then $\{\phi(x, b_i) : i < \omega\}$ is consistent.*

Corollary 1.3. *Suppose T has NIP, μ is a global Keisler measure, and $\phi(x, y) \in L$. Then there do not exist b_i for $i < \omega$ such that the $\mu(\phi(x, b_i)\Delta\phi(x, b_j))$ for $i \neq j$ are bounded away from 0.*

Proof. Suppose there do exist such b_i. We may assume $(b_i : i < \omega)$ is indiscernible (why?) By Lemma 1.2, $\{\phi(x, b_i)\Delta\phi(x, b_{i+1}) : i = 0, 2, 4, ..\}$ is consistent, contradicting NIP.

Corollary 1.3 yields the following important result of Keisler (which we will not be using in these notes).

Corollary 1.4. *Let μ be a Keisler measure over A. Then there is some $B \supset A$ and an extension of μ over B, such that λ has a unique extension to a Keisler measure over any C containing B.*

2 Forking

Basic properties of forking

Definition 2.1.
- Let μ be a global Keisler measure. We say that μ is definable, Borel definable (over A) respectively, if for each $\phi(x, y) \in L$ and closed $C \subset [0, 1]$, $\{b : \mu(\phi(x, b)) \in C\}$ is type-definable (over A), Borel (over A), respectively.

- The global Keisler measure is finitely satisfiable in A if whenever $\mu(X) > 0$ then $X \cap A \neq \emptyset$.
- Suppose $A \subseteq B$ and μ is a Keisler measure over B. If $\phi(x, b)$ (over b) does not divide (fork) over A whenever $\mu(\phi(x, b)) > 0$ we say that μ does not divide (fork) over A.

Lemma 2.2. *(i) Let μ be a global Keisler measure, and A a small set. Then μ divides over A iff μ forks over A.*
(ii) If the global Keisler measure μ is $Autf(\bar{M}/A)$-invariant then μ does not fork over A.
(iii) If μ is either (Borel) definable over A, or finitely satisfiable in A, then μ is $Aut(\bar{M}/A)$-invariant.

Proof.

- (i) If μ forks over A there is $\phi(x)$ with $\mu(\phi(x)) > 0$ and $\models \phi(x) \to \theta_1(x) \vee ... \vee \theta_n(x)$ such that each θ_i divides over A.
- But by finite additivity of μ, some θ_i has positive μ measure, so μ diivides over A.
- (ii) Let $\mu(\phi(x, b)) = r > 0$ and $(b = b_0, b_1, ...)$ an infinite A-indiscernible sequence.
- So $\mu(\phi(x, b_i)) = r$ for all i, so apply Lemma 1.2.

- (iii). This is clear if μ is (Borel) definable over A. Suppose μ is finitely satisfiable in A. Suppose that $tp(b_1/A) = tp(b_2/A)$, and $\phi(x, y) \in L$.
- So $\phi(x, b_1)\Delta\phi(x, b_2)$ is not satisfied in A, whereby $\mu(\phi(x, b_1)\Delta\phi(x, b_2)) = 0$.
- Thus $\mu(\phi(x, b_1)) = \mu(\phi(x, b_2))$ and μ is $Aut(\bar{M}/A)$-invariant.

NIP and forking

We first generalize Theorem 1.6 appropriately.

Theorem 2.3. *(NIP) Let μ be a global Keisler measure. Then the following are equivalent:*

- *(i) μ does not fork over A.*
- *(ii) μ is $Autf(\bar{M}/A)$-invariant.*
- *For any $\phi(x, y) \in L$, whenever $Lstp(b_1/A) = Lstp(b_2/A)$ then $\mu(\phi(x, b_1)\Delta\phi(x, b_2)) = 0$.*

- We first prove (i) implies (iii).
- So suppose μ does not fork over A, and $Lstp(b_1/A) = Lstp(b_2/A)$.
- We can assume that b_1, b_2 begin an A-indiscernible sequence $(b_i : i < \omega)$.
- So $(b_1b_2, b_3b_4, ...)$ is also A-indiscernible.
- If $\mu(\phi(x,b_1)\Delta\phi(x,b_2)) > 0$ then (as μ does not divide over A), we have that $\{\phi(x,b_i)\Delta\phi(x,b_{i+1}) : i = 1,3,...\}$ is consistent, contradicting NIP.
- (iii) implies (i) is obvious, and (ii) implies (i) was in Lemma 2.2.
- This completes the proof of Theorem 2.3.

Theorem II.2.5 also generalizes.

Theorem 2.4. *(NIP) Suppose the global Keisler measure μ does not fork over A. Then μ is $Aut(\bar{M}/bdd^{heq}(A))$-invariant*

Proof.

- We prove that if $tp(b/bdd(A)) = tp(c/bdd(A))$ then $\mu(\phi(x,b)\Delta(\phi(x,c)) = 0$.
- Suppose not. Then $\phi(x,b)\Delta\phi(x,c)$ extends to an ultrafilter in the Boolean algebra of positive μ-measure definable sets. (Explain.)
- This ultrafilter will be precisely a global complete type $p(x)$ which contains $\phi(x,b)\Delta(\phi(x,c))$ and contains no μ-measure 0-formula.
- But then p does not fork over A.
- By Theorem II.2.5, p is $Aut(\bar{M}/bdd^{heq}(A))$-invariant, whence $\phi(x,b) \in p$ iff $\phi(x,c) \in p$, a contradiction.

Averaging

Here some new phenomena enter the picture; averaging a collection of types to obtain an invariant measure.

Theorem 2.5. (NIP) *Let $p(x) \in S(A)$. Then the following are equivalent:*
(i) p does not fork over A (i.e. p has an extension to a global type which does not divide over A).
(ii) p extends to a global Keisler measure μ which is $Aut(\bar{M}/A)$-invariant

- First some remarks.
- The key point is that we have $Aut(\bar{M}/A)$ rather than $Aut(\bar{M}/bdd^{heq}(A))$ in (ii).
- Because any global nonforking extension of p is already $Aut(\bar{M}/bdd^{heq}(A))$-invariant.

- Now for the proof of Theorem 2.5. (ii) implies (i) is immediate, for if μ is as given by (ii) then by Lemma 2.2 μ will not fork over A, so any formula in p will not fork over A.
- (i) implies (ii). We will construct μ and leave verification that it satisfies the required conditions to the reader. Let $\phi(x, y) \in L$, $b \in \bar{M}$ and we want to define $\mu(\phi(x, b))$.
- Let p' be some global nonforking extension of p, which by Corollary II.2.6 is (strongly) Borel definable over $bdd^{heq}(A)$.
- We now discuss a few compact spaces and groups.
- First we have the compact Lascar group or KP-group $G = Aut(bdd^{heq}(A)/A)$, a compact group with its unique (left and right) invariant Haar measure h.
- Second let $S = S_y(bdd^{heq}(A))$ be the space of complete types in variable y over $bdd^{heq}(A)$.
- Let $q(y) = tp(b/A)$ and let $Q \subset S$ be the set of complete extensions of q over $bdd^{heq}(A)$, a closed subspace of S.
- Both S and Q are acted on continuously by G. However Q is also acted on by transitively by G, i.e. is a homogeneous space for G, so has a unique induced G-invariant Borel probability measure h_Q.
- Definability of p' over $bdd^{heq}(A)$ says precisely that the subset X of S consisting of $tp(b'/bdd^{heq}(A))$ such that $\phi(x, b') \in p'$ is a Borel subset of S.
- Hence $X \cap Q$ is a Borel, so measurable, subset of Q.
- Define $\mu(\phi(x, b)) = h_Q(X \cap Q)$.

Uniqueness

- A natural question around Theorem 2.5 is whether there is a unique $Aut(\bar{M}/A)$-invariant global Keisler measure extending $p(x) \in S(A)$ (assuming p does not fork over A).
- If T is stable this will be the case, via the finite equivalence relation theorem.
- Likewise if p (or a global nonforking extension of it) is generically stable (as in Alex U.'s talks).
- But there do exist examples of uniqueness even without generic stability.
- I will formulate a "domination" condition equivalent to uniqueness, and which can be seen as in a sense a measure-theoretic weakening of the statement of the finite equivalence relation theorem

- We assume T has NIP (maybe not necessary) and fix $p(x) \in S(A)$ which has a global nonforking extension.
- Let $\mathcal{P}$ be the set of global nonforking extensions of p. $\mathcal{P}$ is a closed subspace of $S_x(\bar{M})$, so is compact.
- We will identify the $Aut(\bar{M}/A)$-invariant global Keisler measures extending p with the $Aut(\bar{M}/A)$-invariant (regular) Borel probability measures on the space $\mathcal{P}$.
- Let C be the set of extensions of $p(x)$ to complete types over $bdd^{heq}(A)$. C is also a compact space, as well a homogeneous space for the compact Lascar group $G = Aut(bdd^{heq}(A)/A)$, and let h be the unique G-invariant measure on C.
- Let $\pi : \mathcal{P} \to C$ be the canonical (continuous) surjection, taking p to $p|bdd^{heq}(A)$.
- For X a definable (in $\bar{M}$) set of the appropriate sort, let $[X]$ denote the corresponding clopen subset of $\mathcal{P}$.

Definition 2.6. We say that $\mathcal{P}$ is dominated by (C, h, π) if for each definable set X, $\{c \in C : \pi^{-1}(c) \cap [X] \neq \emptyset \text{ and } \pi^{-1}(c) \setminus [X] \neq \emptyset\}$ has h-measure 0.

The following is left to the reader:

Lemma 2.7. *Assume $p(x) \in S(A)$ has a global nonforking extension. Then there is a unique global $Aut(\bar{M}/A)$-invariant Keisler measure extending p if and only if $\mathcal{P}$ is dominated by (C, h, π).*

Example. Consider the real field R together with a new sort X and a regular action of say $SO_2(R)$ on X. Work in a saturated model $\bar{M}$ of this situation. There is a unique type $p(x)$ over set extending "$x \in X$". Moreover there is a unique $Aut(\bar{M})$-invariant global Keisler measure extending p. This is more or less the same example as in the introduction.

VC theorem and Borel definability

We assume that T has NIP. We say a few words (with even fewer proofs) about the relation between the Vapnis-Chervonenkis theorem and the Borel definability of Keisler measures. The main result is

Theorem 2.8. *Suppose μ is a global Keisler measure which does not fork over A. Then μ is Borel definable over $bdd(A)$.*

The main preliminary lemma, which is of interest in its own right, and is a consequence of the VC theorem (probability version) is:

Lemma 2.9. *Let M be any model (even the big model $\bar{M}$ and let μ be a Keisler measure over μ. Let $\phi(x, y) \in L$, and $\epsilon > 0$. Then there are $p_1, .., p_n \in S_x(M)$, such that for any $c \in M$, the difference between $\mu(\phi(x, c))$ and the proportion of p_i which contain $\phi(x, c)$ is $< \epsilon$, and moreover for each $i = 1, .., n$ if $\phi(x, c) \in p_i$ $(\neg\phi(x, c) \in p_i)$, then $\mu(\phi(x, c)) > 0$ $(\mu(\neg\phi(x, c)) > 0)$.*

Proof of Theorem 2.8 from Lemma 2.9.

- Fix $\phi(x, y)$ and $\epsilon > 0$. Let $p_1, .., p_{n_\epsilon}$ be as given by Lemma 2.9.
- As μ does not fork over A, each $p_i|\phi$ does not fork over A so extends to a global complete type $q_i(x)$ which does not fork over A.
- By Theorem II.1.9, each q_i is strongly Borel definable over A.

- Thus for each $i = 1, .., n_\epsilon$ there is a set Y_i^ϵ, a finite Boolean combination of type-definable over $bdd(A)$ sets, such that for any $c \in \bar{M}$, $\phi(x, c) \in p_i$ iff $c \in Y_i$.
- Hence by Lemma 2.8, if c, c' lie in exactly the same Y_i's then $\mu(\phi(x, c))$ and $\mu(\phi(x, c'))$ differ by less than 2ϵ.
- Deducing Borel definability of μ over A from the previous item is left to the reader.
- But for example given a real number r between 0 and 1, to Borel define $\{c' : \mu(\phi(x, c')) = r\}$: we may assume that there is c such that $\mu(\phi(x, c)) = r$.
- The condition on c' is that c' is in precisely those Y_i^ϵ which c is in (as ϵ ranges over $\{1/m : m = 1, 2, ...\}$).