



Ordre et stabilité dans les théories NIP

Pierre Simon

► To cite this version:

Pierre Simon. Ordre et stabilité dans les théories NIP. Mathématiques générales [math.GM]. Université Paris Sud - Paris XI, 2011. Français. NNT : 2011PA112221 . tel-00635463

HAL Id: tel-00635463

<https://theses.hal.science/tel-00635463>

Submitted on 25 Oct 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ PARIS-SUD
FACULTÉ DES SCIENCES D'ORSAY

Thèse présentée pour obtenir le grade de
Docteur en sciences de l'université Paris-Sud

Spécialité : Mathématiques

Ordre et stabilité dans les théories NIP

Pierre SIMON

Soutenue le 21 octobre 2011, devant le jury composé de :

Itai BEN YAACOV
Jean-Benoît BOST
Élisabeth BOUSCAREN (Directrice)
Ehud HRUSHOVSKI
Michael C. LASKOWSKI (Rapporteur)
Anand PILLAY

Rapporteur non présent à la soutenance :

Thomas SCANLON

À mes parents

Remerciements

Je voudrais tout d'abord remercier ma directrice de thèse Élisabeth Bouscaren pour son soutien constant et pour les heures qu'elle a passées à relire mes travaux. J'espère qu'elle me pardonnera si aucun corps différentiel valué n'apparaît dans cette thèse.

Je remercie Udi Hrushovski pour de nombreuses discussions très enrichissantes et pour ses encouragements qui m'ont été précieux. Anand Pillay aussi a partagé beaucoup d'idées avec moi et je lui en suis très reconnaissant. Ce fut un plaisir de collaborer avec Udi et Anand, ainsi qu'avec Artyom Chernikov dont l'optimisme mathématique m'est toujours agréable. Je leur sais gré à tous les trois d'avoir accepté que nos articles écrits en commun soient inclus ici.

Un grand merci aux deux rapporteurs de ma thèse Chris Laskowski et Tom Scanlon pour l'attention qu'ils ont portée à mon manuscrit.

Merci à Itaï Ben Yaacov et à Jean-Benoît Bost d'avoir accepté de faire partie du jury.

Enfin, merci à Saharon Shelah d'avoir créé de si belles mathématiques sans lesquelles ce travail n'existerait pas.

Table des matières

Introduction	vii
0.0.1 Historique et aperçu du sujet	vii
0.0.2 Survol	xi
0.1 Ensembles extérieurement définissables et paires	xii
0.2 Autour des mesures génériquement stables	xiv
0.2.1 L'article <i>On NIP and invariant measures</i> de Hrushovski et Pillay	xiv
0.2.2 Mesures génériquement stables	xv
0.2.3 Constructions de mesures génériquement stables	xvi
0.3 Théories distales	xvii
0.4 Théories dp-minimales ordonnées	xix
1 Préliminaires : Théories NIP	1
1.0.1 VC dimension	1
1.0.2 Formules NIP	3
1.0.3 Types invariants	4
1.0.4 Types génériquement stables	5
1.1 Combinatoire des suites indiscernables	7
1.1.1 Découpage de suites indiscernables	7
1.1.2 Suites mutuellement indiscernables	8
1.1.3 Théories fortement-dépendantes et dp-minimales	8
1.2 Mesures	9
2 Ensembles extérieurement définissables	13
2.1 Introduction	13
2.2 Externally definable sets and honest definitions	15
2.3 On dependent pairs	19
2.4 Applications	22

2.4.1	Naming an indiscernible sequence	22
2.4.2	Dense pairs and related structures	24
3	Mesures génériquement stables	27
3.1	Introduction and preliminaries	27
3.2	Smooth measures and indiscernibles	34
3.3	Generically stable measures	38
3.4	The fsg property for groups, types and measures	43
3.5	Generic compact domination	46
3.6	Borel measures over standard models	52
4	Construction de mesures	55
4.1	Introduction	55
4.2	Preliminaries	56
4.3	Amalgams	60
4.4	Symmetrizations	63
4.5	Smooth measures	67
4.6	σ -additive measures	70
5	Théories distales	73
5.1	Introduction	73
5.1.1	Preliminaries	76
5.1.2	Shrinking	79
5.1.3	Expanding	80
5.1.4	Sliding	81
5.1.5	Weight and dp-minimality	82
5.2	Distal theories	83
5.2.1	Indiscernible sequences	83
5.2.2	Invariant types	87
5.2.3	Generically stable measures	88
5.2.4	Reduction to dimension 1	91
5.3	Domination in non-distal theories	94
5.3.1	Domination for types	97
5.3.2	The finite-co-finite theorem and application	101
5.4	Sharp theories	105
5.4.1	Reduction to dimension 1	108
6	Théories dp-minimales ordonnées	111
6.1	Introduction	111
6.2	Preliminaries on dp-minimality	112
6.3	Inp-minimal ordered structures	114
6.4	Dp-minimal groups	117
6.5	Examples of dp-minimal theories	119

Introduction

Ce travail est une contribution à la théorie des modèles pures. La théorie des modèles est une branche de la logique mathématique dont l'objet est l'étude des *structures* à travers leurs algèbres de Boole des *ensembles définissables*. Une structure est la donnée d'un ensemble M et d'une collection Σ de relations sur M , c'est-à-dire de sous-ensembles de différentes puissances cartésiennes de M . On suppose que Σ contient le graphe de l'égalité. Un ensemble définissable est un élément de $\bigcup_{k < \omega} \mathcal{P}(M^k)$ obtenu en effectuant un nombre fini de combinaisons booléennes, projections, produits et permutations de coordonnées, à partir des relations de Σ .

Ce domaine se décompose traditionnellement en théorie des modèles pures (étude des propriétés combinatoires de la classe des ensembles définissables) et théorie des modèles appliquée (étude des structures particulières, venant de l'algèbre ou de la géométrie). Les deux parties se nourrissent fortement l'une de l'autre : l'étude de structures classiques motive le développement d'une théorie abstraite qui est ensuite appliquée en retour à des cas concrets. Les applications principales se trouvent en géométrie algébrique, en arithmétique et en géométrie réelle.

0.0.1 Historique et aperçu du sujet

Stabilité

On situe habituellement la naissance de la théorie des modèles pure moderne en 1965 avec le théorème de Morley.

Théorème 0.0.1 (Morley, [42]). *Soit T une théorie du premier ordre dans un langage dénombrable. Supposons que pour un cardinal $\kappa > \aleph_0$, T admette, à isomorphisme près, un unique modèle de cardinalité κ , alors ceci est vrai pour tout cardinal $\kappa > \aleph_0$.*

Ce résultat a été le point de départ d'un travail colossal mené par Saharon Shelah, essentiellement dans les années 70 et 80. Shelah s'est posé le problème suivant : étant donnée une théorie complète T , sous quelles conditions les modèles de T sont-ils classifiables, à isomorphismes près, par une collection d'invariants cardinaux ? Par exemple, un corps algébriquement clos de caractéristique nulle est entièrement déterminé par son

degré de transcendance, un \mathbb{Q} -espace vectoriel par sa dimension. À l'inverse, il n'existe pas de collection raisonnable de cardinaux qui permette de caractériser un ordre linéaire à isomorphisme près. Shelah a mis successivement en évidence un certain nombre de propriétés qui doivent être satisfaites par une théorie *classifiable*. Lorsque toutes ces propriétés sont vérifiées par une théorie T donnée, on peut montrer que les modèles de T sont construits autour d'un squelette ayant la forme d'un arbre, un peu de la même manière qu'un espace vectoriel est construit autour d'une de ses bases. Les squelettes possibles sont aisément classifiables, ce qui permet de répondre à la question initiale.

La première de ces propriétés mise en évidence est la plus fondamentale. Il s'agit de la *stabilité*.

Définition 0.0.1. Soit M une structure. Une formule $\phi(\bar{x}, \bar{y})$ est *stable* (dans M) s'il existe un entier N pour lequel il n'existe pas d'uplets $\bar{x}_1, \dots, \bar{x}_N, \bar{y}_1, \dots, \bar{y}_N \in M$ tels qu'on ait :

$$\phi(\bar{x}_i, \bar{y}_j) \iff i \leq j.$$

La structure M est *stable* si toutes les formules y sont stables. Une théorie T est stable si tous ses modèles sont stables.

Une structure infinie qui admet un ordre définissable est instable, puisqu'il suffit de prendre pour $\phi(\bar{x}, \bar{y})$ la formule définissant l'ordre, mais la réciproque est fausse. Plus précisément, Shelah a montré l'alternative suivante : Si T est une théorie instable, alors soit T admet un ordre partiel définissable avec des chaînes infinies, soit une formule $\phi(\bar{x}, \bar{y})$ a la *propriété d'indépendance*, telle que définie ci-dessous.

Définition 0.0.2. Soit M une structure. Une formule $\phi(\bar{x}, \bar{y})$ a la *propriété d'indépendance* (dans M) si pour tout entier N , il existe des uplets $\bar{x}_0, \dots, \bar{x}_{N-1}$ et $(\bar{y}_j : j \in \mathfrak{P}(N))$ d'éléments de M tels qu'on ait :

$$\phi(\bar{x}_i, \bar{y}_j) \iff i \in j.$$

Nous reviendrons plus tard sur cette dernière définition, qui est centrale dans ce travail.

Deux propriétés importantes font la force des théories stables : la définissabilité des types et l'existence d'une bonne notion d'indépendance.

La définissabilité des types est le phénomène suivant : prenons M une structure et $M \prec N$ une extension élémentaire. Soit $\bar{a} \in N$ un uplet et $\phi(\bar{x}, \bar{y})$ une formule. Considérons l'ensemble $\phi(M, \bar{a}) := \{\bar{x} \in M : N \models \phi(\bar{x}, \bar{a})\}$. Cet ensemble est donc la trace sur M d'un ensemble définissable de N et n'est pas en général définissable dans M . On dit que *le type de \bar{a} sur M est définissable* si, pour toute telle formule $\phi(\bar{x}, \bar{y})$, l'ensemble $\phi(M, \bar{a})$ est un ensemble définissable de M . On peut montrer qu'une théorie T est stable si et seulement si tous les types sur tous ses modèles sont définissables.

La deuxième propriété est plus difficile à énoncer précisément. Considérons K un corps algébriquement clos et $C \subseteq K$. On a alors la définition naturelle suivante : deux uplets \bar{a} et \bar{b} sont *indépendants* sur C si $L(\bar{a})$ est algébriquement indépendant de $L(\bar{b})$

au-dessus de L , où L est la clôture algébrique du corps engendré par C . Shelah a défini dans toute théorie stable une notion analogue d'indépendance appelée *non-déviante* et a montré que cette relation satisfait des propriétés naturelles de symétrie et de transitivité. Grâce à cet outil, on peut associer un ensemble de dimensions à une structure stable donnée et finalement, sous certaines hypothèses supplémentaires, classer les modèles de sa théorie.

On dispose ainsi d'un arsenal de techniques pour étudier les structures stables. Malheureusement ces dernières sont assez rares. Ainsi toute structure infinie qui est naturellement ordonnée est instable (telle le corps \mathbb{R} des nombres réels ou tout corps non trivialement valué). Citons comme exemples de structures stables les corps algébriquement clos, les corps séparablement clos, les groupes abéliens, les corps différentiellement clos (l'analogue des corps algébriquement clos pour les corps munis d'une différentielle).

Théories NIP

Une fois que les théories stables ont été bien comprises, les théoriciens des modèles se sont efforcés d'étendre ces idées à des contextes plus larges. En 1992 Hrushovski a montré que les corps pseudofinis admettaient une bonne notion d'indépendance, bien que ces structures soient instables [30] (les corps pseudofinis canoniques sont les ultraproducts non triviaux de corps finis). Motivés par ces résultats, Kim et Pillay ont étudié la notion de théorie *simple* introduite par Shelah 15 ans auparavant. Il s'agit d'une hypothèse combinatoire, plus faible que la stabilité, dont nous ne donnerons pas la définition ici. Ils ont montré ([38]) que cette hypothèse était suffisante pour assurer la présence d'une relation d'indépendance satisfaisant des propriétés naturelles. Les corps pseudofinis sont simples ainsi que les corps algébriquement clos munis d'un automorphisme générique ([13]).

En revanche, aucune structure avec un ordre définissable infini n'est simple. Par conséquent, ni les structures o-minimales (tels les corps réels clos) ni les corps valués n'entrent dans cette classe. Pour intégrer ces structures dans un cadre abstrait, il faut donc étendre la stabilité dans une autre direction.

Nous avons défini ci-dessus la propriété d'indépendance. Elle exprime en quelque sorte qu'on peut coder des ensembles finis par des éléments. Par exemple la formule $x|y$ (x divise y) en arithmétique a la propriété d'indépendance ; il suffit pour le voir de prendre des nombres premiers distincts pour les x_i . Un ensemble de nombres premiers est représenté par leur produit.

On est maintenant en mesure de définir la classe de théories qui est l'objet de ce travail.

Définition 0.0.3. Une théorie T est *NIP* (*No Independence Property*), ou *dépendante*, si dans aucun modèle de T , une formule n'a la propriété d'indépendance.

Il n'est pas difficile de voir qu'une théorie stable est NIP. D'autre part, une théorie simple et NIP est stable. On peut donc considérer NIP comme une extension de la stabilité dans une direction orthogonale à celle de la simplicité. De fait, on perd dans

les théories NIP l'existence d'une bonne notion d'indépendance mais on conserve des succédanés de définissabilité des types (comme on le verra dans le chapitre 1 de ce texte).

Les exemples canoniques de structures NIP sont, en plus des structures stables, les structures o-minimales, les corps valués algébriquement clos ainsi que les corps valués henséliens dont le corps résiduel est stable de caractéristique nulle, et les corps p-adiques. Ce sont donc en quelque sorte des structures construites en combinant des composants stables et des ordres linéaires.

On peut distinguer trois axes ou motivations dans l'étude des théories NIP :

1. Comprendre les propriétés caractéristiques de cette classe. En particulier identifier les théorèmes connus sur les théories stables et o-minimales qui peuvent se généraliser.
2. Dresser une cartographie de l'univers des théories NIP. C'est-à-dire comprendre d'une part les pathologies qui peuvent apparaître et d'autre part identifier des sous-classes pertinentes.
3. Avoir ainsi un ensemble d'outils pour étudier des cas particuliers, par exemple les théories de corps valués et des notions affaiblies d'o-minimalité.

Des travaux importants ont déjà été faits sur chacun de ces points. En ce qui concerne les deux premiers, les résultats principaux sont dus à Shelah. Dans une série d'articles de [55] à [57] il explore les théories NIP dans de nombreuses directions et démontre plusieurs théorèmes majeurs. La diversité des résultats rend difficile d'en rendre compte ici. Nous en mentionnerons cependant certains dans la suite de cette introduction. Une ligne directrice qui semble se dégager de ces travaux est l'intuition suivante (énoncé explicitement à plusieurs reprises) : «Les propriétés des théories NIP s'obtiennent par combinaison de celles des théories stables d'une part et des ordres linéaires d'autre part».

La théorie abstraite a été appliquée essentiellement sur deux sujets : la conjecture de Pillay sur les groupes dans les théories o-minimales (dont nous parlerons plus loin) et l'étude des corps valués. Citons en particulier le travail [31] de Hrushovski et Loeser dans lequel est donnée une construction modèle-théorique naturelle des espaces de Berkovich. Un espace de Berkovich est un certain espace topologique associé à une variété sur un corps valué K . Contrairement à l'espace naturel des K -points de la variété muni de la topologie de la valuation, qui est totalement discontinu, cet espace est localement connexe par arcs. Hrushovski et Loeser ont montré, en utilisant la théorie des modèles des corps valués algébriquement clos et en exploitant le fait que cette théorie est NIP, que cet espace correspondait essentiellement à l'espace des *types génériquement stables* de la variété (dont on donne la définition en 1.0.20. Contentons-nous ici de dire qu'il s'agit de types se comportant comme dans une théorie stable).

On voit là un exemple d'un thème fondamental que nous avons déjà mentionné et qui servira de fil directeur à travers une partie de ce travail : la recherche de composants stables à l'intérieur d'une théorie NIP.

Les théories NIP vues d'ailleurs

La (non-)propriété d'indépendance apparaît dans d'autres domaines des mathématiques (probabilités, géométrie discrète, théorie de l'apprentissage) sous le nom de *VC-dimension finie*. (La définition est donnée dans les préliminaires). Le résultat fondamental à ce sujet est le théorème de Vapnik-Chervonenkis qui est une loi uniforme des grands nombre pour une famille d'événements de VC-dimension finie ([68]). Ce résultat est d'ailleurs utilisé par Hrushovski et Pillay dans l'article [33] en lien avec les mesures de Keisler (voir le paragraphe 0.2.1).

En combinatoire, les familles de VC-dimension finie semblent avoir des propriétés similaires aux familles de convexes de \mathbb{R}^n . Ainsi un certain nombre de résultats connus pour ces derniers (théorème de Helly fractionnaire, théorème (p, q)) ont été démontrés dans le cas de VC-dimension finie ([41]).

Citons enfin une conjecture provenant de la théorie de l'apprentissage dont Laskowski a donné une formulation modèle-théorique essentiellement équivalente ([34]). Il s'agit de la conjecture de définissabilité uniforme des types sur les ensembles finis : si T est une théorie NIP et $\phi(\bar{x}, \bar{y})$ une formule, alors il existe une formule $\psi(\bar{x}, \bar{z})$ telle que pour tout modèle M de T , tout ensemble fini $A \subset M$ et tout $\bar{b} \in M$, il existe $\bar{c} \in A$ tel que $\phi(A, \bar{b}) = \psi(A, \bar{c})$. Elle a été démontrée dans le cas particulier des théories dp-minimales par Guingona, [23].

0.0.2 Survol

Le travail que nous présentons porte sur quatre thèmes différents. Le premier est celui des ensembles extérieurement définissables. La définissabilité des types est un fait fondamental dans les structures stables. Shelah a montré ([56]) que dans une théorie NIP, l'expansion d'un modèle par des définitions pour tous les types, était encore NIP. Ce résultat important est resté mal compris tant à cause du caractère combinatoire de sa preuve que par l'absence d'un cadre plus large dans lequel le placer. Dans le premier chapitre de ce travail, «Ensembles extérieurement définissables et paires de structures NIP» (en commun avec Artem Chernikov), nous définissons une notion d'ensemble *faiblement stablement-plongé* donnant un substitut dans le cas NIP à la définissabilité des types. On en déduit une nouvelle démonstration du résultat de Shelah.

On applique ensuite cette idée au problème des paires de structures NIP. Un certain nombre d'articles ont paru ces dernières années qui étudiaient des structures NIP auxquelles on ajoute un prédicat pour un sous-ensemble particulier. Il peut s'agir soit d'exemples précis (tel \mathbb{C} avec un prédicat pour les racines de l'unité : [7], [11]), soit de situations plus générales (tel une structure topologique avec un prédicat pour une sous-structure dense [9]). Souvent, la structure obtenue reste NIP et cela est montré à l'aide d'un critère ad-hoc et différent à chaque fois. On donne ici des conditions suffisantes pour que l'expansion d'une structure NIP reste NIP qui généralisent presque tous les cas connus.

Le deuxième thème est celui des mesures génériquement stables. Il s'agit d'un objet

nouveau qui existait en filigrane dans l'article [33] de Hrushovski et Pillay, que nous définissons ici et étudions systématiquement. Apparues dans l'étude des groupes définissablement compacts dans les théories o-minimales, ces mesures font un lien intéressant entre théories NIP d'une part et théorie de la stabilité, logique continue et probabilités d'autre part.

Leur étude débute dans le chapitre 2 du présent texte : «Mesures génériquement stables» (travail en commun avec Udi Hrushovski et Anand Pillay) où sont données des définitions équivalentes et propriétés de base. Dans le chapitre 3 : «Constructions de mesures génériquement stables», on donne des constructions permettant d'en produire, simplifiant et généralisant certains résultats précédents.

Le troisième thème s'inscrit dans une problématique que nous avons déjà mentionnée. Il s'agit de comprendre l'existence de comportements stables au sein d'une théorie NIP. Dans le chapitre 3 : «Théories distales», on prend appui sur les résultats précédents pour donner une définition de théorie *purement instable* et on donne des outils pour appréhender la partie stable des types. On donne aussi une application à un résultat partiel d'uniformité lié au premier thème de ce travail, pour lequel aucune approche n'était connue jusqu'alors.

Enfin, le dernier chapitre «Théories dp-minimales ordonnées» traite d'un thème plus particulier que les précédents. Il s'agit d'étudier les théories dp-minimales, qui sont dans un certain sens des théories NIP de *dimension 1*. On se concentre sur les structures dp-minimales ordonnées qui constituent une généralisation assez large des structures o-minimales. L'étude des affaiblissements de la notion d'o-minimalité a été développée dans plusieurs travaux ces dernières années (par ex. [40], [6], [20]). La dp-minimalité est un tel affaiblissement, sans doute le moins géométrique et le plus abstrait de tous ceux considérés jusqu'à présent. On résout notamment deux questions qui se posaient : les groupes ordonnés dp-minimaux sont abéliens et les ensembles définissables dans un tel groupe divisible sont finis ou d'intérieur non vide.

§0.1 Ensembles extérieurement définissables et paires

Ce chapitre, écrit en commun avec Artem Chernikov, a été soumis à l'été 2010 au *Israel Journal of Mathematics* sous le titre «Externally definable sets and dependent pairs».

Il est bien connu qu'une théorie T est stable si et seulement si tout sous-ensemble $A \subseteq \mathcal{C}$ est *stablement plongé*, ce qui veut dire :

(SP) pour toute formule $\phi(x; a)$, avec $a \in \mathcal{C}$, il existe une autre formule $\psi(x, b)$ avec $b \in A$ telle que $\phi(A, a) = \psi(A, b)$.

Quelques observations :

(i) Le cas extrême d'un ensemble non-stablement plongé est le suivant : prenons un graphe biparti aléatoire $M = (U(x), V(y); R(x, y))$. Considérons un petit sous-ensemble $A \subset U$. Alors tout sous-ensemble de A peut s'écrire $R(A, a)$ pour un $a \in V$. Par contre,

les seuls sous-ensembles de A à être relativement définissables par une formule à paramètres dans A sont les ensembles finis ou co-finis. Remarquons que la formule $R(x, y)$ a la propriété d'indépendance.

(ii) Prenons maintenant $A = M$ une structure o-minimale. Les ensembles de la forme $\phi(A, a)$ avec $a \in \mathfrak{C}$, en dimension 1, sont exactement les unions finies d'ensembles convexes. Il y en a donc plus que d'ensembles définissables (M n'est pas stable), mais on voit intuitivement qu'ils ne sont pas trop compliqués et assez proches d'ensembles définissables.

Ceci suggère donc qu'une forme plus faible de (SP) doit être satisfaite par les théories NIP. Plusieurs résultats avaient déjà été démontrés en ce sens. Le plus remarquable est sans doute le théorème suivant de Shelah, généralisant un résultat de Baisalov et Poizat dans le cas o-minimal :

Théorème 0.1.1 (Shelah). *Soit M modèle d'une théorie NIP éliminant les quantificateurs. Notons M^{Sh} la structure obtenue en ajoutant à M un prédicat $R_\phi(\bar{x})$ pour chaque formule $\phi(\bar{x})$ à paramètres dans \mathfrak{C} interprété par $\phi(M)$. Alors M^{Sh} élimine les quantificateurs et est NIP.*

Ce résultat est néanmoins frustrant pour plusieurs raisons : la preuve n'éclaire par beaucoup ce qui se passe et le résultat ne dit rien pour un ensemble A quelconque. De plus (mais nous devons attendre le chapitre 4 pour apporter des précisions à ce sujet), il n'y a pas d'énoncé d'uniformité dans l'élimination des quantificateurs. Pillay a donné une preuve plus simple dans [48], mais sans résoudre ces problèmes.

Le théorème principal que nous démontrons dans le chapitre 1 donne une notion d'ensemble 'faiblement stablement plongé' pertinente dans les théories NIP. Le théorème de Shelah s'en déduit facilement et apparaît ainsi comme une conséquence d'un phénomène plus général.

Théorème 0.1.2. *Soit M modèle d'une théorie NIP et $A \subset M$ un sous-ensemble quelconque. Soit $\phi(x, m)$ une formule à paramètres dans M , alors il existe une extension $(M, A) \prec (N, B)$ de la paire et une formule $\psi(x, b)$ à paramètres dans B telle que $\phi(A, m) = \psi(A, b)$ et que de plus $\psi(B, b) \subseteq \phi(B, m)$.*

La propriété peut-être surprenante (dite d'*honnêteté*) $\psi(B, b) \subseteq \phi(B, m)$ est le point-clé de cet énoncé.

Ces dernières années, un certain nombre de structures construites à partir d'une structure NIP (ou stable) classique en ajoutant un prédicat pour nommer un sous-ensemble particulier ont été étudiées. Une des questions posées à leur sujet est de savoir si la propriété NIP est conservée par cette expansion. En général, cela est vrai. Pour le montrer, il faut toujours établir un critère abstrait qui assure qu'une expansion d'une théorie NIP reste NIP. Plusieurs critères ont ainsi été énoncés, mais tous contiennent des hypothèses de minimalité sur la structure (voir par exemple [11]). On démontre ici un théorème de ce type qui ne suppose pas une telle hypothèse et qui s'applique dans tous les cas connus.

On fixe quelques notations pour pouvoir l'énoncer. On dispose d'un langage L et d'un prédicat unaire additionnel P . On désigne par L_P le langage $L \cup \{P\}$. Soit M une L structure et $A \subset M$ un sous-ensemble. On considère la paire (M, A) comme une L_P -structure d'univers M dans laquelle P est interprété par A .

Une formule est dite *bornée* si elle est de la forme

$$(Q_1 x_1 \in P)(Q_2 x_2 \in P) \dots (Q_n x_n \in P) \phi(x_1, \dots, x_n, y_1, \dots, y_m)$$

où Q_i peut être soit \exists soit \forall et ϕ est une L -formule. La structure $A_{\text{ind}(L)}$ est obtenue en prenant A comme univers et en mettant un prédicat $R_\phi(\bar{x})$ pour chaque L -formule $\phi(\bar{x})$ à paramètres dans M , interprété par la trace de cette formule sur A .

Théorème 0.1.3. *Dans la situation décrite ci-dessus, supposons :*

- (i) M est NIP ;
 - (ii) toute L_P -formule est équivalente dans (M, A) à une formule bornée ;
 - (iii) la structure $A_{\text{ind}(L)}$ est NIP.
- Alors la paire (M, A) est NIP.

Le Théorème 0.1.2 est utilisé pour descendre les paramètres d'une formule donnée à l'intérieur du prédicat P et utiliser ainsi l'hypothèse sur la structure induite sur celui-ci.

§0.2 Autour des mesures génériquement stables

0.2.1 L'article *On NIP and invariant measures* de Hrushovski et Pillay

L'article [33] *On NIP and invariant measures* de Hrushovski et Pillay est le point de départ de la partie centrale de ce travail. Cet article fait suite à [32] écrit avec Peterzil qui traite de la conjecture de Pillay sur les groupes définissablement compacts dans les théories o-minimales. Si T est une théorie o-minimale, un groupe définissable G (affine pour simplifier) est dit *définissablement compact* s'il est fermé et borné. La conjecture de Pillay stipule que si G^{00} désigne le plus petit sous-groupe de G type-définissable d'indice borné, alors G/G^{00} est un groupe de Lie de même dimension que G . Pour démontrer cette conjecture, les auteurs utilisent un outil introduit par Keisler pour étudier les théories NIP dans [37] mais qui n'avait alors pas encore trouvé d'applications : les mesures sur l'espace des types.

Une mesure de Keisler est une mesure finiment additive sur l'algèbre de Boole des ensembles définissables. De manière équivalente, c'est une mesure borélienne régulière sur l'espace des types. On peut aussi voir les mesures comme des types au sens de la logique continue, mais c'est un point de vue que nous n'adopterons pas ici afin de rester dans le cadre de la logique du premier ordre. Les théories NIP sont un cadre naturel pour étudier les mesures car elles y jouissent de bonnes propriétés, les rendant manipulables essentiellement comme des types (par exemple, leur support dans l'espace des types est un ensemble de cardinalité bornée).

Si G est un groupe définissablement compact, Hrushovski et Pillay montrent dans [33] qu'il admet une mesure de Keisler invariante par translation et que cette mesure est unique. L'image de cette mesure dans le groupe de Lie G/G^{00} n'est autre que la mesure de Haar sur ce groupe. La mesure invariante donne donc une mesure nulle à tous les ensembles infinitésimaux. Pour montrer cela, les auteurs prouvent qu'un tel groupe satisfait une hypothèse technique appelée *fsg* (*génériques finiment satisfaisables*). Ils montrent ensuite que dans une théorie NIP, un groupe *fsg* admet une unique mesure de Keisler invariante.

De cela est déduit un théorème dit *de domination compacte* de G par le groupe de Lie G/G^{00} . Notons π la projection canonique de G sur G/G^{00} . Alors si $D \subseteq G$ est un ensemble définissable, l'ensemble des $g \in G/G^{00}$ tels que $\pi^{-1}(g)$ a une intersection non vide avec D et son complémentaire a mesure de Haar nulle. Ainsi les ensembles définissables sont bien approximés en un certain sens par leur image dans le quotient. Le point clé de la preuve est l'unicité de la mesure invariante sur G .

Enfin, les auteurs étudient la notion de *type génériquement stable* introduite précédemment par Shelah. Un type est génériquement stable s'il est définissable et finiment satisfaisable sur un ensemble borné. De manière équivalente, c'est un type global invariant dont la suite de Morley est totalement indiscernable. Un cas particulier de tel type est donné par les types stablement dominés étudiés dans [27] et repris ensuite dans [31] pour donner une construction modèle-théorique des espaces de Berkovich que nous avons déjà mentionnée.

La question est posée dans cet article de la bonne généralisation de cette notion dans le cadre des mesures. On répond à cela dans le chapitre 2 écrit en commun avec Hrushovski et Pillay.

0.2.2 Mesures génériquement stables

Ce chapitre, écrit en commun avec Ehud Hrushovski et Anand Pillay a été accepté pour publication au *Transactions of the AMS* sous le titre «On generically stable and smooth measures in NIP theories».

Dans cet article, on généralise aux mesures la définition de type génériquement stable. On y montre que les différentes définitions équivalentes ont un analogue naturel dans ce cadre ce qui prouve le bienfondé de cette notion. En particulier, les mesures génériquement stables sont définissables et finiment satisfaisables. Les groupes *fsg* apparaissent comme les groupes admettant une mesure invariante génériquement stable. L'unicité de la mesure invariante sur un tel groupe, démontrée dans [33], est alors un corollaire de la théorie.

Une construction standard en théorie des modèles permet, étant donné un groupe définissable G , de construire une structure dans laquelle ce même groupe G apparaît comme un (sous-groupe d'un) groupe d'automorphismes. (Voir l'introduction de [33], essentiellement on construit un espace homogène sur G). Il en découle une certaine analogie entre groupes définissables et groupes d'automorphismes. Dans cette optique,

des analogues de la propriété de fsg et de domination compacte dans le cas de groupes d'automorphismes sont conjecturés dans [33]. Ainsi, une définition de fsg pour un type sur un ensemble A quelconque est proposée et les auteurs demandent si un tel type p s'étend de manière unique en une mesure A -invariante. On montre ici que la définition suggérée est équivalente à l'existence d'une mesure génériquement stable A -invariante étendant p . On peut alors répondre positivement à la question.

On généralise aussi la notion de domination compacte au cas de la domination par un espace de types. Plus précisément, considérons un ensemble de paramètres A satisfaisant une hypothèse de clôture ($A = \text{bdd}(A)$, réalisée par exemple si A est un modèle). Une mesure μ sur A admet une unique extension en une mesure globale A -invariante si et seulement si on a la propriété suivante de domination :

☒ Pour tout ensemble \mathfrak{C} -définissable D , non-déviant sur A , l'ensemble des types sur A admettant à la fois une extension A -invariante satisfaisant p et une extension satisfaisant $\neg D$, est de μ -mesure nulle.

Enfin, on fait deux observations (la première étant beaucoup plus simple que la seconde) :

(i) Si $\langle a_i : i \in (0, 1) \rangle$ est une suite indiscernable, on définit une mesure μ sur \mathfrak{C} par $\mu(\phi(x)) = \lambda_0(\{t \in (0, 1) : \models \phi(a_t)\})$, où λ_0 désigne la mesure de Lebesgue. Alors μ est génériquement stable. Ceci permet de construire facilement des mesures génériquement stables.

(ii) Une mesure σ -additive sur \mathbb{R} ou \mathbb{Q}_p induit une mesure de Keisler lisse (c'est-à-dire ayant une unique extension à un modèle plus grand ; une propriété beaucoup plus forte que la générique stabilité, que nous considérons comme l'analogue pour les mesures d'un type réalisé). Ce théorème généralise des résultats de Karpinski et Macintyre [35] concernant la définissabilité des mesures de Haar réelle ou p -adique sur le disque unité.

On voit donc que les mesures génériquement stables sont plus universelles que les types du même nom. Une thèse qu'on développe un peu plus dans le chapitre suivant.

0.2.3 Constructions de mesures génériquement stables

Ce chapitre a été accepté pour publication au *Journal of Symbolic Logic* sous le titre «Finding generically stable measures».

Le but de ce chapitre est de montrer que les mesures génériquement stables sont courantes et se comportent mieux que les types génériquement stables. On y généralise les deux observations (i) et (ii) ci-dessus.

D'abord, en partant de n'importe quelle mesure μ , on peut construire une mesure μ_Σ génériquement stable appelée une symétrisation de μ . L'idée est simple : de même que le type moyen d'une suite totalement indiscernable d'éléments est génériquement stable, on peut construire une mesure génériquement stable en moyennant une suite totalement indiscernable de mesures. Or on peut facilement construire une telle suite à partir de la suite de Morley $\mu_{x_1, x_2, \dots}^{(\omega)}$ de μ en prenant la moyenne sur les permutations des variables

x_1, x_2, \dots . Ceci permet par exemple de démontrer simplement qu’une mesure génériquement stable sur une sorte imaginaire est la projection d’une mesure génériquement stable sur M (prendre un relèvement quelconque et le symétriser). L’analogue est faux pour les types génériquement stables. Ensuite, on montre que toute mesure σ -additive sur un modèle fixé muni d’une tribu suffisamment riche induit une mesure génériquement stable :

Théorème 0.2.1. *Soit M une structure NIP et \mathcal{B} une tribu de parties de M . On suppose que pour toute extension $M \prec N$ et toute formule $\phi(x_1, x_2; a)$ avec $a \in N$, la trace $\phi(M^2; a)$ de cette formule est mesurable pour la tribu produit $\mathcal{B}^{\otimes 2}$. Soit λ une mesure de probabilité sur (M, \mathcal{B}) , alors λ induit une mesure de Keisler génériquement stable sur M .*

Enfin, pour retrouver complètement les résultats de [29] sur les mesures σ -additives, on montre que toute mesure génériquement stable sur une structure o-minimale ou élémentairement équivalente à \mathbb{Q}_p est lisse.

On voit là apparaître une propriété des corps o-minimaux et p-adiques qu’on pourrait qualifier de ‘pure instabilité’ ; elle implique par exemple qu’il n’existe pas de type génériquement stable non trivial. Il est naturel à ce point de se demander si on peut caractériser cette propriété sans faire appel aux mesures, ou si elle découle d’une autre propriété plus forte. Par ailleurs, la preuve présentée dans ce chapitre fonctionne par récurrence sur la dimension en utilisant les fonctions de Skolem définissables et laisse ouverte la question générale : si toutes les mesures génériquement stables sont lisses en dimension 1, en est-il de même en toutes dimensions ?

On répond à ces questions dans le chapitre suivant.

§0.3 Théories distales

Ce chapitre a été soumis pour publication aux *Annals of Pure and Applied Logic* en juin 2011 sous le titre «Distal and non-distal theories».

Cet article est composée de deux parties essentiellement indépendantes. La première répond aux questions soulevées dans le Chapitre 3 sur les théories NIP où toutes les mesures génériquement stables sont lisses. Ces théories sont appelées ‘distales’. La deuxième partie étudie les théories non-distales et montre que dans un certain sens, les relations non distales entre deux types sont des relations stables.

a) Théories distales.

Motivé par l’exemple des corps valués, on considère souvent les théories NIP comme étant une combinaison de parties stables et d’ordres. Dans le cas des corps algébriquement clos valués, cette idée est rendue précise par les travaux de Haskell, Hrushovski et Macpherson sur la métastabilité ([27]) : dans cette théorie, un type $\text{tp}(a/C)$ (sur une ‘bonne base’ C) a une composante o-minimale $\Gamma(a)$ et un quotient $\text{tp}(a/\Gamma(a)C)$ génériquement stable (et même stablement dominé, ce qui est une condition plus forte). On ne

peut probablement pas s'attendre à ce qu'un résultat similaire soit vrai en général. On prend néanmoins ici cette intuition à la lettre et on cherche à définir dans une théorie NIP quelconque une *partie stable*.

Bien sûr, on ne sait pas exactement ce qu'on entend par là. L'idée de départ est la suivante : au lieu de définir ce qu'est la partie stable d'une théorie, on peut chercher ce qui caractérise les théories pour lesquelles cette partie stable est triviale. On pense en particulier aux théories o-minimales ou aux corps p-adiques. On veut donc définir une notion de théorie NIP purement instable.

Une condition évidemment nécessaire pour qu'une théorie soit qualifiée de purement instable est qu'elle n'admette pas d'ensemble interprétable infini stable (c'est-à-dire sur lequel la structure induite est stable). Mais cela semble bien trop faible. Une définition plus forte serait :

(D0) Il n'y a pas de types génériquement stables (non-réalisés).

Cette définition a un problème majeur : elle n'est pas conservée par passage de M à M^{eq} . Une solution serait de la formuler directement sur M^{eq} . Seulement, il est probable que dans le cas NIP, M^{eq} ne contienne pas assez d'éléments et qu'il devienne nécessaire d'introduire des imaginaires d'ordre supérieur (comme cela a été le cas pour les théories simples).

Au chapitre précédent il est apparu une version plus forte de (D0) :

(D1) Les mesures génériquement stables sont lisses.

On sait que cette notion est stable par passage de M à M^{eq} et qu'elle est satisfaite par les théories o-minimales et les corps p-adiques. Les résultats de cette première partie montrent que cette définition est robuste et se comporte mieux que (D0). C'est donc un bon candidat pour la définition recherchée. Dans ce chapitre, on donne des définitions équivalentes en terme de suites indiscernables et de types invariants et on prouve qu'il suffit de vérifier (D1) en dimension 1. Plus précisément :

Théorème 0.3.1. *Il y a équivalence, pour une théorie NIP T , entre les trois propriétés suivantes :*

(i) *Pour toute suite indiscernable $I_1 + I_2$, les types limites $\lim(I_1)$ et $\lim(I_2)$ sont orthogonaux,*

(ii) *Deux types invariants qui commutent sont orthogonaux,*

(iii) *Toute mesure génériquement stable est lisse.*

De plus, il suffit de vérifier chacune de ces conditions en dimension 1.

Une théorie satisfaisant (D1) est appelée «distale».

b) Fort de cette définition, on se propose dans une deuxième partie d'étudier les théories NIP non-distales, suivant l'intuition que la non-distalité doit correspondre à des phénomènes de stabilité. Le résultat principal est ainsi la mise en évidence d'une relation d'indépendance $\mathbf{a} \downarrow_M^s \mathbf{b}$ qui intuitivement dit que les parties stables de \mathbf{a} et \mathbf{b} sont indépendantes au-dessus de M . On montre que cette relation est symétrique et de poids borné.

Ainsi, on n'a pas exhibé explicitement la partie stable, mais on a réussi à donner un sens à la phrase «les parties stables de \mathbf{a} et \mathbf{b} sont indépendantes au-dessus de \mathbf{M} ».

On obtient comme application un théorème sur la combinatoire des suites indiscernables ('Finite-co-finite theorem' 5.3.30), dont nous ne connaissons pas de démonstration directe (malgré son apparente simplicité).

Théorème 0.3.2. *Soit $I = I_1 + I_2 + I_3$ une suite indiscernable avec I_1 et I_2 infinis, et soit $\phi(x, \mathbf{a})$ une formule. On suppose que $I_1 + I_3$ est une suite indiscernable sur les paramètres \mathbf{a} . Alors l'ensemble $\{x \in I_2 : \models \phi(x, \mathbf{a})\}$ est fini ou co-fini dans I_2 .*

c) Dans la dernière partie de ce chapitre, on définit une classe de théories (les théories *sharp*, ou *nettes* en français) pour lesquelles la 'partie stable' est matérialisée par des types génériquement stables. Pour justifier cette définition, on donne un critère en terme de décomposition des suites indiscernables et on montre qu'il suffit de vérifier ce critère sur les suites d'éléments (et non d'uplets). En particulier, toute théorie dp-minimale est nette (voir paragraphe suivant).

§0.4 Théories dp-minimales ordonnées

Ce chapitre a été publié au *Journal of Symbolic Logic*, 76, 2 (2011), pp. 448-460, sous le titre «On dp-minimal ordered structures».

La question de savoir s'il existe une notion abstraite de minimalité incluant les structures o-minimales est restée ouverte pendant longtemps. Shelah dans [53] a proposé la *dp-minimalité* qui est une forme uni-dimensionnelle de la propriété NIP (voir préliminaires 1.1.2). Voici quelques exemples de théories dp-minimales :

1. Les théories superstables de rang U égal à 1,
2. Les théories o-minimales,
3. Les corps valués algébriquement clos,
5. Plus généralement les structures C-minimales (qui sont aux corps valués algébriquement clos ce que les structures o-minimales sont aux corps réels clos),
4. Les p-adiques.

Les théories dp-minimales ordonnées ont donc été étudiées très récemment comme une généralisation assez large des structures o-minimales. Ainsi, les groupes dp-minimaux ordonnés ont été étudiés par Goodrick dans [21] où il prouve qu'une fonction définissable unaire est union d'un nombre fini de fonctions continues localement monotones. Cette recherche s'inscrit dans un mouvement important de tentatives d'affaiblissement de la notion d'o-minimalité. De nombreuses définitions ont ainsi été données (structures faiblement o-minimales, quasi o-minimales, *o-minimal open core*... Voir par exemple [39], [6], [20]) dont certaines seulement impliquent la dp-minimalité.

Dans cette optique, on s'intéresse surtout aux structures qui sont des expansions d'un groupe ordonné, voire d'un corps réel clos. Néanmoins, la particularité de l'étude qu'on fait ici est de ne pas introduire de loi de groupe tout de suite et d'essayer de tirer

le maximum d'information de la seule structure d'ordre. On montre ainsi qu'un type dévie sur un modèle M si et seulement si sa réduction à l'ordre pur dévie. Dans un deuxième temps seulement, on s'intéresse aux groupes ordonnés. On résout alors deux questions qui se posaient à leur sujet : les groupes ordonnés dp-minimaux sont abéliens, et les ensembles définissables en dimension 1 sont finis ou d'intérieur non-vide (dans le cas d'un groupe divisible). D'autre part, on donne une démonstration simple du fait qu'un groupe dp-minimal est *abélien-par-exposant-fini*. Ceci était connu pour les groupes C-minimaux ([39]).

Enfin, on donne des exemples de structures dp-minimales. On montre que les ordres totaux et les arbres purs sont dp-minimaux (il est connu depuis longtemps que ces structures sont NIP). Ceci est à rapprocher du fait que les structures o-minimales ou C-minimales sont NIP : un ordre pur n'est pas nécessairement o-minimal et un arbre pur peut ne pas être C-minimal, mais ce résultat montre qu'ils sont tout de même relativement simples.

1.0.1 VC dimension

Nous commençons par introduire la notion fondamentale de VC-dimension. Nous prenons dans cette section le point de vue combinatoire, mais nous l'abandonnerons rapidement pour adopter un point de vue plus qualitatif et modèle-théorique.

On considère un ensemble E et une famille \mathcal{C} de parties de E .

Définition 1.0.1. Des éléments x_1, \dots, x_n sont *éclatés* par \mathcal{C} si pour toute partie $I \subseteq \{1, \dots, n\}$, il existe $A \in \mathcal{C}$ tel que : $x_i \in A \Leftrightarrow i \in I$.

Définition 1.0.2. La classe \mathcal{C} est de VC-dimension n s'il existe n points éclatés par \mathcal{C} , mais pas $n + 1$.

Elle est dite de VC-dimension infinie s'il existe n points éclatés par \mathcal{C} pour tout n .

Quelques exemples : Si (E, \leq) est un ensemble totalement ordonné, et \mathcal{C} est la collection des intervalles, alors $\text{VC-dim}(\mathcal{C})=2$.

Soit (E, \leq) un arbre (c'est-à-dire que \leq définit un ordre partiel et pour tout $a \in E$, l'ensemble $\{x \in E : x \leq a\}$ est totalement ordonné par \leq). On prend pour \mathcal{C} l'ensemble des boules fermées : $\mathcal{C} = \{F_a : a \in E\}$ où $F_a = \{x \in E : x \geq a\}$. Alors on a $\text{VC-dim}(\mathcal{C})=2$.

Ces deux exemples sont les exemples fondamentaux de classes de VC-dimension finie. Ils sont reliés en théorie des modèles respectivement aux structures o-minimales et C-minimales.

Avant de parler de théorie des modèles, citons deux résultats combinatoires. Le premier est instructif, mais ne sera pas utilisé par la suite. Le deuxième est le théorème fondamental de Vapnik et Chervonenkis qui a justifié l'étude des familles de VC-dimension finie en géométrie combinatoire et en théorie de l'apprentissage. Il n'interviendra pas directement dans cette thèse, mais est sous-jacent à certains résultats sur les mesures.

Lemme de Sauer

Étant donnée une famille \mathcal{C} comme ci-dessus, on définit

$$\pi_{\mathcal{C}}(\mathfrak{m}) = \max_{B \subseteq A, |B|=\mathfrak{m}} |\{S \cap B : S \in \mathcal{C}\}|.$$

Lemme 1.0.3 (Lemme de Sauer). *Si la famille \mathcal{C} est de VC-dimension au plus d , alors $\mathcal{C}(\mathfrak{m}) = O(\mathfrak{m}^d)$.*

Ainsi, pour une famille \mathcal{C} donnée, soit $\pi_{\mathcal{C}}(\mathfrak{m})$ vaut $2^{\mathfrak{m}}$ pour tout \mathfrak{m} , soit cette fonction est à croissance polynomiale.

Théorème de Vapnik-Chervonenkis et ϵ -réseaux

Nous décrivons maintenant le résultat fondateur de la notion de VC-dimension.

Soit (X, Ω, μ) un espace de probabilités (Ω est une tribu sur l'ensemble X et μ une mesure probabilité sur cette tribu). Soit aussi \mathcal{C} une famille de sous-ensembles de X . Pour chaque entier k , on note μ^k la mesure produit sur l'espace X^k .

On considère les hypothèses suivantes :

- (i) Chaque $C \in \mathcal{C}$ est μ -mesurable,
- (ii)_d La famille \mathcal{C} est de VC-dimension au plus d ,

Si $C \subseteq X$, on note $\text{Fr}(C; x_1, \dots, x_n) = \frac{1}{n} \sum_i \chi_C(x_i)$ où χ_C est la fonction caractéristique de C . On définit aussi, pour chaque entier n ,

$$g_{\mathcal{C}}^k(x_1, \dots, x_k) = \sup_{C \in \mathcal{C}} |\text{Fr}(C; x_1, \dots, x_k) - \mu(C)|,$$

et

$$h_{\mathcal{C}}^k(x_1, \dots, x_k, y_1, \dots, y_k) = \sup_{C \in \mathcal{C}} |\text{Fr}(C, x_1, \dots, x_k) - \text{Fr}(C, y_1, \dots, y_k)|.$$

(iii) Pour chaque k , la fonction $g_{\mathcal{C}}^k$ est μ^k -mesurable et la fonction $h_{\mathcal{C}}^k$ est μ^{2k} -mesurable.

Théorème 1.0.1 (Vapnik-Chervonenkis). *Soit d un entier, alors il existe une fonction $f_d(k, \epsilon)$ telle que pour tout $\epsilon > 0$, $f_d(k, \epsilon) \rightarrow 0$ lorsque $k \rightarrow +\infty$ et tel que pour tout espace de probabilités (X, Ω, μ) et toute famille \mathcal{C} satisfaisant aux conditions (i), (ii)_d, (iii) ci-dessus, on ait :*

$$\mu^k(\{\bar{x} : g_{\mathcal{C}}^k(\bar{x}) > \epsilon\}) < f(k, \epsilon).$$

Ce théorème est une version uniforme de la loi des grands nombres. Il dit que dans une suite de tirages aléatoires, la fréquence d'apparition d'un certain événement C converge vers la mesure de C , et ce uniformément en $C \in \mathcal{C}$.

Un corollaire immédiat, est l'existence d' ϵ -réseaux pour les familles de VC-dimension finie.

Corollaire 1.0.4. Soit (X, Ω, μ) un espace de probabilités et \mathcal{C} une famille de parties mesurables de X de VC-dimension finie. Soit aussi $\epsilon > 0$. Il existe alors un ensemble fini $A \subset X$ tel que pour tout $C \in \mathcal{C}$, on ait

$$\left| \frac{|A \cap C|}{|A|} - \mu(C) \right| \leq \epsilon.$$

Pour plus de détails, voir [41], Chapitre 10. Voir aussi [10] pour l'importance de ces notions en théorie de l'apprentissage.

1.0.2 Formules NIP

On se place maintenant dans le cadre de la théorie des modèles de la logique du premier ordre. On utilise les notations standard. On travaille ainsi avec une théorie T donnée ; \mathfrak{C} désigne un modèle $\bar{\kappa}$ -saturé et $\bar{\kappa}$ -homogène pour un grand cardinal $\bar{\kappa}$ supérieur à tous les cardinaux des sous-structures considérées. Un sous-ensemble de \mathfrak{C} sera dit *borné* s'il est de cardinalité inférieure à $\bar{\kappa}$.

On désignera par $S(A)$ l'espace des types complets sur A . Si $A \subseteq B$ sont deux ensembles et p un type sur B , on notera $p|_A$ la restriction de p à A .

Nous donnons maintenant la définitions des théories NIP et présentons leurs propriétés principales. Pour plus de détails, nous renvoyons le lecteur à l'article d'exposition de H. Adler [2] et aux sections 1 à 4 de l'article [33] de Hrushovski et Pillay.

Définition 1.0.5. Une formule $\phi(x, y)$ (x et y a priori de tailles différentes) a la *propriété d'indépendance* si la VC-dimension de \mathcal{C} est infinie, où $\mathcal{C} = \{\phi(x, b) \mid b \in \mathfrak{C}\}$.

Dans le cas contraire, on dira que ϕ est *NIP* (*not independence property*).

Remarque 1.0.6. Par compacité, la propriété d'indépendance équivaut à : pour tout λ il existe $\{a_i, i < \lambda\}$ et $\{b_I, I \in \mathcal{P}(\lambda)\}$ tels que :

$$\models \phi(a_i, b_I) \Leftrightarrow i \in I.$$

En particulier, ceci implique qu'il existe 2^λ types sur l'ensemble des a_i , et ainsi que T est instable.

Une théorie est *NIP* si toutes ses formules sont NIP.

Remarque 1.0.7. Le fait d'être NIP est symétrique en x et y , c'est à dire que si $\phi(x, y)$ est NIP, $\psi(y, x) := \phi(x, y)$ l'est aussi.

Si ϕ et ψ sont NIP, il en va de même de $\neg\phi$, $\phi \wedge \psi$ et $\phi \vee \psi$.

Le théorème suivant donne une caractérisation équivalente de la propriété d'indépendance qui est souvent plus simple à exploiter que la définition donnée ci-dessus. Rappelons qu'une suite $(a_i)_{i < \omega}$ est dite *indiscernable* si pour tous $i_1 < i_2 < \dots < i_n$ et $j_1 < j_2 < \dots < j_n$ et toute formule $\phi(x_1, \dots, x_n)$ on a

$$\phi(a_{i_1}, \dots, a_{i_n}) \Leftrightarrow \phi(a_{j_1}, \dots, a_{j_n}).$$

Une suite est dite *totalemtent indiscernable* si ceci reste vrai pour toute permutation des indices.

Proposition 1.0.8. *Une formule $\phi(x, y)$ a la propriété d'indépendance si et seulement s'il existe une suite indiscernable $(a_i)_{i < \omega}$ et un uplet b tel qu'on ait pour tout i :*

$$\models \phi(a_{2i}, b) \wedge \neg \phi(a_{2i+1}, b).$$

Ainsi, à toute formule NIP $\phi(x, y)$ est associé un entier N vérifiant : Pour toute suite indiscernable (a_i) , il n'existe pas de b tel que $\neg(\phi(a_i, b) \leftrightarrow \phi(a_{i+1}, b))$ pour $i = 0, \dots, N$. Le plus petit tel nombre s'appelle le *nombre d'alternances* de ϕ .

Par conséquent, si T est NIP et que $(a_i)_{i \in J}$ (où J est un ordre linéaire sans dernier élément) est une suite indiscernable, pour tout ensemble A de paramètres, la suite

$$(\text{tp}(a_i/A) : i \in J)$$

converge dans $S(A)$. La limite est appelée type limite de (a_i) sur A et notée $\lim((a_i)_{i \in J}/A)$.

Supposons que la suite $(a_i)_{i \in J}$ est de plus totalement indiscernable (c'est-à-dire que toute permutation de cette suite est indiscernable). Soit $\phi(x, y)$ une formule NIP de nombre d'alternance N et $b \in \mathcal{C}$. Supposons que $\lim((a_i)/b) \models \phi(x, b)$. Alors il existe au plus N valeurs de i telles que $a_i \models \neg \phi(x, b)$.

Théorème 1.0.2 (Shelah). *Une théorie est NIP si et seulement si toutes les formules $\phi(x, y)$, où x est une seule variable (et non un uplet de variables), le sont.*

On peut maintenant donner des exemples de théories NIP.

Proposition 1.0.9. *Les théories (ou classes de théories) suivantes sont NIP.*

1. *Les théories stables,*
2. *Les théories o-minimales,*
3. *Les théories C-minimales (analogue de la o-minimalité pour des arbres à la place d'ordres linéaires), par exemple ACVF : la théorie des corps valués algébriquement clos,*
4. *Toute théorie d'un groupe abélien ordonné,*
5. *Une théorie d'un corps valué hensélien de caractéristique résiduelle nulle et de corps résiduel k a la propriété d'indépendance si et seulement si $\text{Th}(k)$ l'a,*
6. *La théorie des corps p -adiques (dans le langage des corps pur ou dans celui des corps valués).*

1.0.3 Types invariants

Si $A \subset \mathcal{C}$ est un ensemble de paramètres, un type global $p \in S(\mathcal{C})$ est dit A -invariant, s'il est invariant par tout automorphisme qui fixe A . En d'autres termes, pour toute formule $\phi(x; y)$ et $b, b' \in \mathcal{C}$, si b et b' ont même type sur A , alors $p \models \phi(x; b) \leftrightarrow \phi(x; b')$. On dira qu'un type global p est invariant s'il est A -invariant pour un A de taille $< \bar{\kappa}$. Un type invariant est entièrement déterminé par sa restriction à un modèle $M \supset A$, $|A|^+$ -saturé.

Il est vrai dans toute théorie qu'un type M -invariant ne dévie pas sur M . Dans une théorie NIP, la réciproque l'est aussi.

Lemme 1.0.10. *Soit M un modèle d'une théorie NIP T et $p \in S(\mathfrak{C})$ un type global. Alors p ne dévie pas sur M si et seulement s'il est M -invariant.*

Soient p_x et q_y deux types globaux invariants. On peut définir le produit $(p \otimes q)_{xy}$. Pour cela, prendre $a \models p|_{\mathfrak{C}}$ et $b \models q|_{\mathfrak{C}a}$. On pose alors $(p \otimes q) = \text{tp}(ab/\mathfrak{C})$. C'est un type invariant. En général, $p_x \otimes q_y \neq q_y \otimes p_x$ (voir le Chapitre 5 pour une étude de ces questions).

Soit p un type global A -invariant et $M \subset \mathfrak{C}$ contenant A . On construit inductivement une suite $(a_i)_{i < \omega}$ de la manière suivante : on prend $a_0 \models p|_M$, puis $a_1 \models p|_{Ma_0}$, $a_2 \models p|_{Ma_0a_1}$ et ainsi de suite. On dit que $(a_i)_{i < \omega}$ est une *suite de Morley* de p sur M . C'est une suite indiscernable sur M dont le type ne dépend que de p et M . Dans le cas $M = \mathfrak{C}$, on notera ce type p^ω .

Dans une théorie stable, on retrouve p à partir de $(a_i)_{i < \omega}$ par la formule $p = \lim((a_i)/\mathfrak{C})$. Ceci n'est plus vrai en général dans une théorie NIP (prendre pour p le type à $+\infty$ dans les ordres linéaires denses). Néanmoins, les types invariants sont quand même déterminés par le type de leur suite de Morley comme l'indique la proposition suivante.

Proposition 1.0.11. *Supposons T NIP. Soient p et q deux types globaux A -invariants. Si $p^{(\omega)}|_A = q^{(\omega)}|_A$, alors $p = q$.*

Nous décrivons maintenant une procédure pour retrouver le type A -invariant p à partir d'une réalisation $(a_i)_{i < \omega}$ de sa suite de Morley sur A . Soit b un paramètre quelconque. Soit a_ω réalisant $\lim((a_i)_{i < \omega}/Ab)$. On cherche, si possible, un élément $a_{\omega+1}$ tel que $(a_i)_{i \leq \omega+1}$ soit indiscernable sur A et $a_\omega \models \phi(x, b) \iff a_{\omega+1} \models \neg\phi(x, b)$. On continue ensuite en cherchant $a_{\omega+2}$ tel que $a_{\omega+1} \models \phi(x, b) \iff a_{\omega+2} \models \neg\phi(x, b)$. Et ainsi de suite. Par NIP, cette construction doit s'arrêter à un certain $a_{\omega+k}$. On a alors $p \vdash \phi(x, b) \iff a_{\omega+k} \models \phi(x, b)$. En effet, on peut construire un $a_{\omega+k+1}$ en prenant une réalisation de p sur tout ce qui précède.

1.0.4 Types génériquement stables

On définit maintenant plusieurs notions de types stables (dans une théorie générale d'abord, puis dans le cas particulier d'une théorie NIP).

Commençons par discuter les notions d'ensembles stables.

Définition 1.0.12. Un sous-ensemble A de \mathfrak{C} est *stablement plongé* si tout sous-ensemble définissable de A est définissable avec paramètres dans A .

(On appelle ensemble définissable de A l'intersection avec A d'un ensemble définissable de \mathfrak{C} .)

De manière équivalente, A est stablement plongé si tous les types sur A sont définissables. On voit donc qu'une théorie est stable si et seulement si tous les ensembles sont stablement plongés.

Dans le cas général, cette notion sera surtout utilisée pour A définissable ou type-définissable.

Proposition 1.0.13. *Soit X un ensemble type-définissable. Les propriétés suivantes sont équivalentes :*

1. X est stablement plongé.
2. Pour tout \mathbf{a} , il existe $X_0 \subseteq X$ borné tel que $\text{tp}(\mathbf{a}/X_0) \vdash \text{tp}(\mathbf{a}/X)$.
3. Pour tout \mathbf{a} , $\text{tp}(\mathbf{a}/X)$ est définissable sur un X_0 borné.
4. Tout automorphisme de X_0 s'étend en un automorphisme de X .

Définition 1.0.14. Un ensemble type-définissable X est *faiblement stable* s'il n'existe pas de formule $\phi(\mathbf{x}, \mathbf{y})$ (à paramètres dans \mathfrak{C}) et des $(\mathbf{a}_i, \mathbf{b}_i)_{i < \omega}$ uplets d'éléments de X tels que $\mathfrak{C} \models \phi(\mathbf{a}_i, \mathbf{b}_j)$ si et seulement si $i \leq j$.

Définition 1.0.15. Un ensemble type-définissable X est *stable*, *stablement plongé* s'il n'existe pas de formule $\phi(\mathbf{x}, \mathbf{y})$ (à paramètres dans \mathfrak{C}) et des $(\mathbf{a}_i, \mathbf{b}_i)_{i < \omega}$, ou les \mathbf{a}_i sont des uplets d'éléments de X tels que $\mathfrak{C} \models \phi(\mathbf{a}_i, \mathbf{b}_j)$ si et seulement si $i \leq j$.

Proposition 1.0.16. *Supposons, pour simplifier, T dénombrable. Soit X type-définissable. Les propriétés suivantes sont équivalentes :*

1. X est stable, stablement plongé.
2. Pour tout $A \subset \mathfrak{C}$ dénombrable et toute formule $\phi(\mathbf{x}, \mathbf{y})$, l'ensemble $\{\text{tp}_\phi(\mathbf{c}/A) \mid \mathbf{c} \in X\}$ est dénombrable.
3. Pour tout A , le nombre de A -types réalisés dans X est inférieur à $|A|^{\aleph_0}$.
4. Pour tout A et $\mathbf{x} \in X$, le type de \mathbf{x} sur A est définissable.

Proposition 1.0.17. *Supposons X \emptyset -définissable, alors X est stable, stablement plongé si et seulement si X est faiblement stable et stablement plongé.*

Dans le cas NIP, les choses sont un peu plus simples.

Proposition 1.0.18. *Supposons T NIP, alors si l'ensemble type-définissable X est faiblement stable, il est stable stablement plongé.*

On définit à présent les types génériquement stables, qui se comportent génériquement comme des types dans une théorie stable.

Proposition 1.0.19. *(NIP) Soit \mathbf{p} un type global A -invariant. Les conditions suivantes sont équivalentes :*

1. \mathbf{p} est définissable et finiment satisfaisable dans un modèle contenant A .
2. Toute suite de Morley de \mathbf{p} sur A est totalement indiscernable.
3. Pour tout $B \supseteq A$ borné, \mathbf{p} est l'unique extension globale non-déviant de $\mathbf{p}|_B$.
4. Si I réalise $\mathbf{p}^{(\omega)}|_A$, alors $\mathbf{p} = \text{Lim}(I/\mathfrak{C})$.

Définition 1.0.20. Un type invariant vérifiant les conditions équivalentes ci-dessus est dit *génériquement stable*.

Plus généralement, un type sur un ensemble A sera dit génériquement stable s'il admet une extension non-déviant génériquement stable. (Ceci ne dépend pas de l'extension choisie).

Il est facile de voir que l'existence d'un type génériquement stable (non-réalisé) est équivalente à celle d'une suite totalement indiscernable (non-constante). La condition 2 donne une implication, et réciproquement, si $I = (a_i)_{i < \omega}$ est totalement indiscernable, alors le type limite $p = \lim(I/\mathfrak{C})$ est un type global I -invariant. Il est clairement finiment satisfaisable dans I et aussi définissable par : $p \vdash \phi(x, c)$ si et seulement si $\models \bigvee_{w \subset 2N, |w|=N} \bigwedge_{i \in w} \phi(a_i, c)$ pour N assez grand.

Enfin, on dit qu'un type est *pleinement stable* si toutes ses extensions (déviantes et non-déviantes) sont génériquement stables.

EXEMPLE 1.0.21. *Tout type contenant une formule stable stablement plongé est pleinement stable. Un exemple d'un tel ensemble est donné par le corps résiduel dans un modèle d'ACVF (corps valués algébriquement clos).*

Un exemple de type génériquement stable non pleinement stable est donné par le type générique d'une boule fermée dans un modèle d'ACVF.

§1.1 Combinatoire des suites indiscernables

On sait qu'une théorie est NIP si et seulement si pour toute suite indiscernable $I = (a_i)_{i \in \mathcal{I}}$ et toute formule $\phi(x, b)$, on ne peut pas trouver de suite strictement croissante $(i_k : k < \omega)$ d'éléments de \mathcal{I} vérifiant $\phi(a_{i_k}, b) \leftrightarrow \neg \phi(a_{i_{k+1}}, b)$ pour tout k . On étudie dans cette section des conséquences et raffinements de cette propriété.

On suppose dans toute cette section que T est une théorie NIP.

1.1.1 Découpage de suites indiscernables

La propriété appelée caractérise les sous-ensembles de \mathcal{I} qui peuvent s'écrire $\{i \in \mathcal{I} : \models \phi(a_i, b)\}$ pour une certaine formule $\phi(x, b)$. On va maintenant donner une caractérisation similaire pour les sous-ensembles de \mathcal{I}^n qui s'expriment comme $\{(i_1, \dots, i_n) \in \mathcal{I}^n : \models \phi(a_{i_1}, \dots, a_{i_n}, b)\}$ pour une certaine formule $\phi(x_1, \dots, x_n, b)$.

Une relation d'équivalence \sim sur \mathcal{I} est dit convexe si ses classes d'équivalences sont convexes. On écrira $(i_1, \dots, i_n) \sim (j_1, \dots, j_n)$ si pour tout $k \leq n$, on a $i_k \sim j_k$ et pour tout $k, k' \leq n$, on a $i_k \leq i_{k'} \iff j_k \leq j_{k'}$.

Enfin, une relation d'équivalence convexe \sim est finie si elle a un nombre fini de classes. Elle sera dite *essentiellement de taille κ* si elle est l'intersection de κ relations d'équivalence convexes finies. Si \sim est essentiellement de taille κ , elle peut avoir jusqu'à 2^κ classes.

Théorème 1.1.1. *Soit $I = (a_i)_{i \in \mathcal{I}}$ une suite indiscernable. Soit b un uplet fini de paramètres et $\phi(x_1, \dots, x_n; y)$ une formule. Alors il existe une relation d'équivalence convexe finie \sim sur \mathcal{I} telle que :*

$$(i_1, \dots, i_n) \sim (j_1, \dots, j_n) \implies \phi(a_{i_1}, \dots, a_{i_n}; b) \leftrightarrow \phi(a_{j_1}, \dots, a_{j_n}; b).$$

On voit que si on prend pour \sim la relation d'équivalence la plus grossière ayant cette propriété, les coupures qu'elle induit sur I sont définissables à partir de \mathbf{b} dans la structure (\mathfrak{C}, I) où on a ajouté un prédicat unaire pour I . On dira de ces coupures qu'elles sont induites par \mathbf{b} sur I .

Si on fait parcourir à ϕ toutes les formules à paramètres dans \mathbf{b} , on obtient $|T|$ coupures de J définies par les différents \sim , et l'intersection de toutes ses relations d'équivalence donne une relation d'équivalence essentiellement de taille $|T|$. Il est remarquable que si on travaille sur un ensemble de paramètres A de taille quelconque, cela n'augmente pas le nombre de coupures induites par \mathbf{b} .

Théorème 1.1.2 (Découpage d'une suite indiscernable). *Soit $A \subset \mathfrak{C}$ un ensemble quelconque de paramètres et soit $I = (a_i)_{i \in J}$ une suite indiscernable sur A . Soit \mathbf{b} un uplet fini. Il existe alors une relation d'équivalence convexe \sim sur J essentiellement de taille $|T|$ telle que pour tout \mathbf{n} et toute formule $\phi(x_1, \dots, x_n; y)$ à paramètres dans A on ait :*

$$(i_1, \dots, i_n) \sim (j_1, \dots, j_n) \implies \phi(a_{i_1}, \dots, a_{i_n}; \mathbf{b}) \leftrightarrow \phi(a_{j_1}, \dots, a_{j_n}; \mathbf{b}).$$

Un cas particulier important de ces résultats (qui peut se prouver directement) est que si I est une suite A -indiscernable indexée par un ordre J de cofinalité $> |T|$, et \mathbf{b} est un uplet fini, alors il existe un segment final de I qui est indiscernable sur $A\mathbf{b}$.

1.1.2 Suites mutuellement indiscernables

Une famille $(I_i)_{i \in \Omega}$ est dite *mutuellement indiscernable* si pour tout $i \in \Omega$, I_i est indiscernable sur $\cup_{j \neq i} I_j$. Par exemple, si $I = \sum_{i \in \Omega} I_i$ est une suite indiscernable (où on note $I + J$ la concaténation des deux suites I et J), alors les suites $(I_i)_{i \in \Omega}$ sont mutuellement indiscernables. De manière générale, les énoncés sur le découpage d'une suite indiscernable admettent un analogue en termes de suites mutuellement indiscernables. Par exemple :

Théorème 1.1.3 (Découpage en termes de suites mutuellement indiscernables). *Soit $A \subset \mathfrak{C}$ et $(I_i)_{i \in \Omega}$ une famille de suites mutuellement indiscernables sur A . Soit \mathbf{d} un uplet fini. Alors il existe un sous ensemble $\Omega' \subseteq \Omega$ avec $|\Omega \setminus \Omega'| \leq |T|$ tel que la famille $(I_i)_{i \in \Omega'}$ soit mutuellement indiscernable sur $A\mathbf{d}$.*

Un exemple de suites mutuellement indiscernables est donné par la situation suivante : soit $(p_i)_{i \in \Omega}$ une famille de types globaux A -invariants. On suppose que les p_i commutent deux-à-deux, c'est-à-dire que pour $i \neq j \in \Omega$, on a $p_i \otimes p_j = p_j \otimes p_i$. On réalise inductivement $I_i \models p_i^{(\omega)} | A \cup \{I_j : j <_\Omega i\}$. Alors la famille $(I_i)_{i \in \Omega}$ est mutuellement indiscernable.

1.1.3 Théories fortement-dépendantes et dp-minimales

Le théorème précédent suggère naturellement des notions de 'forte-dépendance' analogues à la superstabilité. Shelah en a ainsi étudié plusieurs dans [53]. Nous donnons ici la principale.

Définition 1.1.1. Soit T une théorie NIP. Alors T est dite *fortement-dépendante* si pour toute famille $(I_i)_{i \in \Omega}$ de suites mutuellement indiscernables, et tout uplet \mathbf{d} fini, il existe un sous-ensemble $\Omega' \subseteq \Omega$ avec $|\Omega \setminus \Omega'| < \omega$ tel que pour tout $i \in \Omega'$, la suite I_i soit indiscernable sur \mathbf{d} .

Enfin, la notion de théorie dp-minimale est encore plus restrictive.

Définition 1.1.2. Une théorie NIP T est dite *dp-minimale* si pour toute famille $(I_i)_{i \in \Omega}$ de suites mutuellement indiscernables, et tout 1-uplet \mathbf{d} , il existe $i_0 \in \Omega$ tel que pour tout $i \in \Omega$, $i \neq i_0$, I_i est indiscernable sur \mathbf{d} .

Une théorie stable de rang U finie est fortement-dépendante, mais la réciproque n'est pas vraie. Les théories o-minimales, C-minimales et les corps p-adiques sont dp-minimales.

Voir le Chapitre 6 de cette thèse pour plus d'informations sur les théories dp-minimales.

§1.2 Mesures

On introduit à présent les mesures de Keisler. Pour plus de précisions à ce sujet, voir les articles [32] et [33].

On suppose dans toute cette section que T est NIP.

Si A est un ensemble de paramètres, on note $\mathcal{L}_x(A)$ l'ensemble des formules en la variable libre x et à paramètres dans A quotienté par la relation d'équivalence $T \vdash \phi \leftrightarrow \psi$. On identifie cet ensemble avec l'ensemble des sous-ensembles A -définissables de $\mathfrak{C}^{|x|}$.

Définition 1.2.1. Une *mesure de Keisler* (ou *mesure*) en la variable x sur A est une mesure de probabilité finiment additive sur l'algèbre de Boole $\mathcal{L}_x(A)$.

C'est-à-dire qu'on a pour toutes formules $\phi(x)$ et $\psi(x)$,

$$\mu(\phi(x) \wedge \psi(y)) + \mu(\phi(x) \vee \psi(x)) = \mu(\phi(x)) + \mu(\psi(x)),$$

et

$$\mu(x = x) = 1.$$

On notera parfois une mesure μ_x pour indiquer que la mesure μ porte sur la variable x .

EXEMPLE 1.2.2. Un type est un cas particulier de mesure : soit $p \in S_x(A)$, alors p peut être vu comme une mesure en posant, pour $\phi \in \mathcal{L}_x(A)$, $p(\phi) = 1$ si $p \vdash \phi$ et $p(\phi) = 0$ dans le cas contraire.

Pour μ une mesure, l'ensemble $S(\mu) = \{p \in S(A) \mid \forall \phi (\mu(\phi) = 1 \rightarrow p \models \phi)\}$ est un fermé de l'espace des types appelé *support* de μ .

On note $\mathcal{M}_x(A)$ l'espace des mesures de Keisler sur A en la variable x . C'est un espace compact lorsqu'on le munit de la topologie faible, c'est-à-dire celle engendrée par les $B_\phi(x_0, r) := \{\mu : |\mu(\phi) - x_0| < r\}$ pour $x_0, r \in [0, 1]$.

Toute mesure μ sur A s'étend de manière unique en une mesure borélienne régulière sur l'espace des types $S(A)$ («régulière» veut dire que la mesure d'un borélien X est l'infimum des mesures des ouverts O contenant X . Dans notre cas où l'espace est totalement discontinu, la mesure de O est elle-même le supremum des mesures des ouverts-fermés qu'il contient). Réciproquement, toute mesure borélienne régulière sur $S(A)$ donne une mesure de Keisler en la restreignant aux ouverts-fermés.

On a donc une bijection :

$$\text{Mesures de Keisler sur } A \longleftrightarrow \text{Mesures boréliennes régulières sur } S(A).$$

Quelques résultats élémentaires issus de [32] et [33] :

Proposition 1.2.3. *Si (a_i) est une suite indiscernable et $\phi(x, y)$ une formule, supposons qu'il existe $\epsilon > 0$ tel que $\mu(\phi(x, a_i)) > \epsilon$ pour tout i , alors $\bigwedge \phi(x, a_i)$ est consistant.*

Proposition 1.2.4. *Il n'existe pas $\{b_i, i < \omega\}$ et $\epsilon > 0$ vérifiant*

$$i \neq j \rightarrow \mu(\phi(x, b_i) \Delta \phi(x, b_j)) > \epsilon.$$

À partir d'une mesure μ , on peut définir une relation d'équivalence sur les ensembles définissables par $\phi \sim_\mu \psi$ si $\mu(\phi \Delta \psi) = 0$.

Proposition 1.2.5. *L'ensemble $\{X / \sim_\mu\}$ est borné (i.e., de taille $< \bar{\kappa}$).*

Corollaire 1.2.6. *L'espace $S(\mu)$ est de cardinalité bornée.*

EXEMPLE 1.2.7. *Contre-exemple à ces propriétés dans le cas non-NIP : prendre pour T la théorie du graphe aléatoire (sur le langage $L = \{R\}$) et pour μ la mesure globale définie par : $\mu(\bigwedge_{i=1}^n (x R a_i)^{\epsilon_i}) = 2^{-n}$ (pour tout $\epsilon_i \in \{-1, 1\}$ et $a_i \in \mathfrak{C}$). Le support de μ est l'ensemble des types non-réalisés.*

On peut déduire du théorème de Vapnik-Chervonenkis 1.0.1, le résultat suivant :

Proposition 1.2.8. *Soit μ une mesure globale et $\phi(x, y)$ une formule. Soit aussi $\epsilon > 0$. Il existe alors $p_1, \dots, p_n \in S(\mu)$ tels que pour tout $b \in \mathfrak{C}$:*

$$\left| \mu(\phi(x, b)) - \frac{1}{n} \sum_{i \leq n} p_i(\phi(x, b)) \right| \leq \epsilon.$$

Ces résultats expliquent que dans les théories NIP, les mesures ont un comportement proche de celui des types. On peut généraliser les définitions usuelles.

Définition 1.2.9. Soient $M \prec N$, avec N $|M|^+$ -saturé et soit $\mu \in \mathcal{M}(N)$,

- μ est *finiment satisfaisable* dans M si pour tout $\phi \in L_x(N)$ tel que $\mu(\phi) > 0$, il existe $a \in M$ tel que $N \models \phi(a)$.
- μ est *M-invariante* si pour tout $\phi(x; y) \in L$, et $b \equiv_M b'$, $\mu(\phi(x; b)) = \mu(\phi(x; b'))$.

- μ est *définissable* sur M si elle est M -invariante et pour tout $\phi(x; y) \in L$, et $r \in [0, 1]$, l'ensemble $\{p \in S_y(M) : \mu(\phi(x; b)) \leq r \text{ pour tout } b \in N, b \models p\}$ est un fermé de $S_y(M)$.
- μ est Borel-définissable sur M si l'ensemble ci-dessus est un borélien de $S_y(M)$.

Proposition 1.2.10 ([33] 4.9). *Si $\mu \in \mathcal{M}(N)$ est M -invariants (N est $|M|^+$ -saturé), alors μ est Borel-définissable sur M .*

Cette propriété permet de définir le produit de deux mesures $\mu_x \otimes \mu_y$ de la manière suivante. Soit $M \prec N$, N un modèle $|M|^+$ -saturé. Soient $\mu_x \in \mathcal{M}(N)$ une mesure M -invariante et $\lambda_y \in \mathcal{M}(N)$ une mesure quelconque. Alors la mesure $\mu_x \otimes \lambda_y$ est définie comme une mesure à deux variables x, y sur N par $(\mu_x \otimes \lambda_y)(\phi(x, y)) = \int_{S_y(P)} f(y) d\lambda_y$, où $P \prec N$ est un petit modèle contenant M et les paramètres de ϕ et $f : S(P) \rightarrow [0, 1]$ est définie par $f(p) = \mu(\phi(x, b))$ pour $b \in N, b \models p$.

Si μ_x est une mesure globale M -invariante, on définit par récurrence : $\mu_{x_1 \dots x_n}^{(n)}$ par $\mu_{x_1}^{(1)} = \mu$ et $\mu_{x_1 \dots x_{n+1}}^{n+1} = \mu_{x_{n+1}} \otimes \mu_{x_1 \dots x_n}^{(n)}$. On définit $\mu_{x_1 x_2 \dots}^{(\omega)}$ comme l'union des $\mu^{(n)}$. C'est la *suite de Morley* de μ . C'est une suite indiscernable au sens suivant.

Définition 1.2.11. Une mesure $\mu_{x_1 x_2 \dots}$ est indiscernable sur A si pour tout $\phi(x_1, \dots, x_n) \in L(A)$ et des indices $i_1 < \dots < i_n$, on a

$$\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{i_1}, \dots, x_{i_n})).$$

On aura besoin du résultat suivant de [71] (voir aussi Chapitre 2, 2.10).

Proposition 1.2.12. *Si $\mu_{x_1, x_2, \dots} \in \mathcal{M}(M)$ est indiscernable, et $\omega_{y, x_1, x_2, \dots}$ étend μ , alors pour toute formule $\phi(x, y) \in L(M)$, $\lim_{i \rightarrow \omega} \omega(\phi(x_i, y))$ existe. De manière équivalente, pour tout $\phi(x, y)$ et $\epsilon > 0$, il y a un indice N tel que pour toute mesure $\omega_{y, x_1, x_2, \dots}$ comme ci-dessus, on ait $|\omega(\phi(x_i, y)) - \omega(\phi(x_{i+1}, y))| \geq \epsilon$ pour au plus N valeurs de i .*

Mesures lisses

Nous considérons la notion suivante comme un analogue pour les mesures des types réalisés.

Définition 1.2.13 (Lisse). Une mesure $\mu \in \mathcal{M}(N)$ est lisse si μ a une unique extension globale. Si $M \subset N$, on dit que μ est lisse sur M si $\mu|_M$ est lisse.

Les propriétés importantes suivantes sont prouvées dans le Chapitre 2.

Proposition 1.2.14 (II. 2.3). *Soit μ lisse sur M et soient $\phi(x, y) \in L$ et $\epsilon > 0$. Alors il existe des formules $v_i^1(x)$, $v_i^2(x)$ et $\psi_i(y)$ pour $i = 1, \dots, n$ dans $L(M)$ telles que :*

1. Pour tout $b \in \mathfrak{C}$, il existe i tel que $b \models \psi_i(y)$;
2. pour tout i et $b \in \mathfrak{C}$, si $b \models \psi_i(y)$, alors $\mathfrak{C} \models v_i^1(x) \rightarrow \phi(x, b) \rightarrow v_i^2(x)$;
3. pour tout i , $\mu(v_i^2(x)) - \mu(v_i^1(x)) < \epsilon$.

Réciproquement, si la condition est satisfaite pour tout $\phi(x, y)$ et ϵ , alors μ est lisse.

Corollaire 1.2.15. *Si μ est lisse sur N , alors :*

1. *il existe $M \prec N$ de taille $|T|$ tel que μ est lisse sur M ;*
2. *μ est définissable et finiment satisfaisable dans N (en particulier μ est génériquement stable).*

Le fait suivant a été observé dans l'article fondateur de Keisler [37].

Lemme 1.2.16 (II. 2.2). *Soit μ une mesure sur M . Alors il y a une extension $M \prec N$ et une mesure μ' sur N étendant μ telle que μ' est lisse.*

Ensembles extérieurement définissables et paires

écrit en commun avec Artem Chernikov

§2.1 Introduction

This paper is organised in two main parts, the first studies externally definable sets in first order NIP theories and the second, using those results, proves dependence of some theories with a predicate, under quite general hypothesis. We believe both parts to be of independent interest. A third section gives some examples of dependent pairs and relates results proved here to ones existing in the literature.

Honest definitions

Let M be a model of a theory T . An *externally definable* subset of M^k is an $X \subseteq M^k$ that is equal to $\phi(M^k, d)$ for some formula ϕ and d in some $N \succ M$. In a stable theory, by definability of types, any externally definable set coincides with some M -definable set. By contrast, in a random graph for example, any subset in dimension 1 is externally definable.

Assume now that T is NIP. A theorem of Shelah ([53]), generalising a result of Poizat and Baisalov in the o-minimal case ([3]), states that the projection of an externally definable set is again externally definable. His proof does not give any information on the formula defining the projection. A slightly clarified account is given by Pillay in [48].

In section 1, we show how this result follows from a stronger one : existence of honest definitions. An *honest definition* of an externally definable set is a formula $\phi(x, d)$ whose trace on M is X and which implies all M -definable subsets containing X . Then the projection of X can be obtained simply by taking the trace of the projection of $\phi(x, d)$.

Combining this notion with an idea from [22], we can adapt honest definitions to be defined over any subset A instead of a model M . We obtain a property of *weak stable-embeddedness* of sets in NIP structures. Namely, consider a pair (M, A) , where we have

14 CHAPITRE 2. ENSEMBLES EXTÉRIEUREMENT DÉFINISSABLES

added a unary predicate $\mathbf{P}(x)$ for the set A . Take $c \in M$ and $\phi(x, c)$ a formula. We consider $\phi(A, c)$. If A is stably embedded, then this set is A -definable. Guingona shows that in an NIP theory, this set is externally A -definable, *i.e.*, coincides with $\psi(A, d)$ for some $\psi(x, y) \in L$ and $d \in A'$ where $(M', A') \succ (M, A)$. We strengthen this by showing that one can find such a $\phi(x, d)$ with the additional property that $\psi(x, d)$ never lies, namely $(M', A') \models \psi(x, d) \rightarrow \phi(x, c)$. In particular, the projection of $\psi(x, d)$ has the same trace on A as the projection of $\phi(x, c)$. This is the main tool used in Section 2 to prove dependence of pairs.

Dependent pairs

In the second part of the paper we try to understand when dependence of a theory is preserved after naming a new subset by a predicate. We provide a quite general sufficient condition for the dependence of the pair, in terms of the structure induced on the predicate and the restriction of quantification to the named set.

This question was studied for stable theories by a number of people (see [12] and [4] for the most general results). In the last few years there has been a large number of papers proving dependence for some pair-like structures, *e.g.* [9], [25], [11], *etc.* We apologise for adding yet another result to the list. However, our approach differs in an important way from the previous ones, in that we work in a general NIP context and do not make any assumption of minimality of the structure (by asking for example that the algebraic closure controls relations between points). In particular, in the case of pairs of models, we obtain that if M is dependent, $N \succ M$ and (N, M) is bounded (see Section 2 for a definition), then (N, M) is dependent.

Those results seem to apply to most, if not all, of the pairs known to be dependent. It also covers some new cases, in particular answering a question of Baldwin and Benedikt about naming an indiscernible sequence.

The setting

We will not make a blanket assumption that T is NIP, so we work a priori with a general first order theory T in a language L . We use standard notation. We have a monster model \mathbb{M} . If A is a set of parameters, $L(A)$ denotes the formulas of L with parameters from A . If $\phi(x)$ is some formula, and A a subset of \mathbb{M} , we will write $\phi(A)$ for the set of tuples $a \in A^{|x|}$ such that $\phi(a)$ holds. If A is a set of parameters, by $\phi(x) \rightarrow^A \psi(x)$, we mean that for every $a \in A$, $\phi(a) \rightarrow \psi(a)$ holds. Also $\phi(x) \rightarrow^{p(x)} \psi(x)$ stands for $\phi(x) \rightarrow^{p(\mathbb{M})} \psi(x)$.

We will often consider pairs of structures. So if our base language is L , we define the language $L_{\mathbf{P}}$ where we add to L a new unary predicate $\mathbf{P}(x)$. If M is an L -structure and $A \subseteq M$, by the pair (M, A) we mean the $L_{\mathbf{P}}$ extension of M obtained by setting $\mathbf{P}(a) \Leftrightarrow a \in A$. Throughout the paper $\mathbf{P}(x)$ will always denote this extra predicate.

As usual $\text{alt}(\phi)$ is the maximal number n such that there exists an indiscernible sequence $(a_i)_{i < n}$ and c satisfying $\phi(a_i, c) \Leftrightarrow i$ is even. Standardly $\phi(x, y)$ is dependent

if and only if $\text{alt}(\phi)$ is finite. For more on the basics of dependent theories see e.g. [2].

Acknowledgments

We are grateful to Itay Kaplan and the referee for suggesting a number of improvements to the paper, and to Itai Ben Yaacov and Manuel Bleichner for pointing out some typos and deficiencies.

§2.2 Externally definable sets and honest definitions

Recall that a partial type $p(x)$ is said to be *stably embedded* if any definable subset of $p(x)$ is definable with parameters from $p(\mathbb{M})$. It is well known that if $p(x)$ is stable, then $p(x)$ is stably embedded (see e.g. [44]). We are concerned with an analogous property replacing stable by dependent.

We say that a formula $\phi(x, c)$ is NIP over a (partial) type $p(x)$ if there is no indiscernible sequence $(a_i)_{i < \omega}$ of realisations of p such that $\phi(a_i, c)$ holds if and only if i is even. We say that $\phi(x, y)$ is NIP over $p(x)$ if $\phi(x, c)$ is NIP over $p(x)$ for every c .

The following is the fundamental observation. We assume here that we have two languages $L \subseteq L'$, and we work inside a monster model \mathbb{M} that is an L' -structure. The language L' could be L_p for example.

Proposition 2.2.1. *Let $p(x)$ be a partial L' -type and $\phi(x, c) \in L(\mathbb{M})$ be NIP over $p(x)$. Then for each small $A \subseteq p(\mathbb{M})$ there is $\theta(x) \in L(p(\mathbb{M}))$ such that*

- 1) $\theta(x) \cap A = \phi(x, c) \cap A$
- 2) $\theta(x) \rightarrow^{p(x)} \phi(x, c)$
- 3) $\phi(x, c) \setminus \theta(x)$ does not contain any A -invariant global L -type consistent with $p(x)$.

Démonstration. Let $q(x) \in S_L(\mathbb{M})$ be A -invariant and consistent with $\{\phi(x, c)\} \cup p(x)$. We try to choose inductively $a_i, b_i \in p(\mathbb{M})$ and $q_i \subseteq q$, for $i < \omega$ such that

- $q_i(x) = q(x)|_{A_{a_i < b_i}}$
- $a_i \models q_i(x) \cup \{\phi(x, c)\} \cup p(x)$ (we can always find one by assumption)
- $b_i \models q_i(x) \cup \{\neg\phi(x, c)\} \cup p(x)$.

Assume we succeed. Consider the sequence $(d_i)_{i < \omega}$ where $d_i = a_i$ if i is even and $d_i = b_i$ otherwise. It is a Morley sequence of q over A , and as such is L -indiscernible. Furthermore, we have $\models \phi(d_i, c)$ if and only if i is even. This contradicts $\phi(x, y)$ being NIP over $p(x)$, so the construction must stop at some finite stage i_0 . Then $q_{i_0}(x) \rightarrow^{p(x)} \phi(x, c)$ and by compactness there is $\psi_q(x) \in q_{i_0}$ (so $\psi_q \in L(p(\mathbb{M}))$) such that $\psi_q(x) \rightarrow^{p(x)} \phi(x, c)$. So we see that the set of all such ψ_q 's covers the compact space of global L -types invariant over A and consistent with $\{\phi(x, c)\} \cup p(x)$ (so in particular all realised types of elements of A such that $\phi(a, c)$). Let $(\psi_j)_{j < n}$ be a finite subcovering, then taking $\theta(x) = \bigvee_{j < n} \psi_j(x)$ does the job. \square

Definition 2.2.2 (Externally definable set). Let M be a model, an externally definable set of M is a subset X of M^k for some k such that there is a formula $\phi(x, y)$ and $d \in M$ with $\phi(M, d) = X$. Such a $\phi(x, d)$ is called a definition of X .

16 CHAPITRE 2. ENSEMBLES EXTÉRIEUREMENT DÉFINISSABLES

We can now prove a form of *weak stable embeddedness* for NIP formulas.

Corollary 2.2.3 (Weak stable-embeddedness). *Let $\phi(x, y)$ be NIP. Given (M, A) and $c \in M$ there are $(M', A') \succeq (M, A)$ and $\theta(x) \in L(A')$ such that $\phi(A, c) = \theta(A)$ and $\theta(x) \rightarrow^{A'} \phi(x, c)$.*

Démonstration. Notice that $\phi(x, y)$ is still NIP in any expansion of the structure. In particular in the L_P -structure (M, A) . Now apply Proposition 2.2.1 with $L' = L_P$ and $p(x) = \{P(x)\}$. \square

Question 2.2.4. *Do we get uniform weak stable embeddedness? In other words, is it possible to choose θ depending just on ϕ , or at least just on ϕ and $\text{Th}(M, A)$?*

Corollary 2.2.5. *Let $f : M \rightarrow M$ be an externally definable function, that is the trace on M of an externally definable relation which happens to be a function on M . Then there is an \mathbb{M} -definable partial function $g : \mathbb{M} \rightarrow \mathbb{M}$ with $g|_M = f$.*

Démonstration. Let $\phi(x, y; c)$ induce f on M , $c \in N \succ M$. By Corollary 2.2.3 we find $(N', M') \succ (N, M)$ and $\theta(x, y) \in L(M')$ satisfying $\theta(M^2) = \phi(M^2, c)$ and $\theta(x, y) \rightarrow^{M'} \phi(x, y; c)$. As the extension of pairs is elementary and $M' \models T$, it follows that $\theta(x, y)$ is a graph of a partial function. \square

Definition 2.2.6 (Honest definition). Let $X \subseteq M^k$ be externally definable. Then an honest definition of X is a definition $\phi(x, d)$ of X , $d \in \mathbb{M}$ such that :

$$\mathbb{M} \models \phi(x, d) \rightarrow \psi(x) \text{ for every } \psi(x) \in L(M) \text{ such that } X \subseteq \psi(M).$$

In Section 2, we will need the notion of an honest definition *over* A which is defined at the beginning of that section.

Proposition 2.2.7. *Let T be NIP. Then every externally definable set $X \subset M^k$ has an honest definition.*

Démonstration. Let $M \prec N$ and $\phi(x) \in L(N)$ be a definition of X , and let $(N', M') \succeq (N, M)$ be $|N|^+$ -saturated (in L_P). Let $\theta(x) \in L(M')$ as given by Corollary 2.2.3, so $(N', M') \models (\forall x \in P) \theta(x) \rightarrow \phi(x)$. If $\psi(x) \in L(M)$ with $X \subseteq \psi(M)$ then $(N', M') \models (\forall x \in P) \phi(x) \rightarrow \psi(x)$. Combining, we get $(N', M') \models (\forall x \in P) \theta(x) \rightarrow \psi(x)$. But since $M' \models T$ and $\theta(x), \psi(x) \in L(M')$ we have finally $M' \models \theta(x) \rightarrow \psi(x)$. \square

We illustrate this notion with an o-minimal example inspired by [3].

We let M_0 be the real closure of \mathbb{Q} and let $\epsilon > 0$ be an infinitesimal element. Let M be the real closure of $M_0(\epsilon)$. Let π be the usual transcendental number, and finally let N be the real closure of $M(\pi)$.

Lemma 2.2.8. *Let $0 < b \in N$ be infinitesimal, then there is $n \in \mathbb{N}$ such that $b < \epsilon^{1/n}$.*

Démonstration. We define a valuation v on $\mathbb{Q}(\pi, \epsilon)$ by setting $v(x) = 0$ for all $x \in \mathbb{Q}(\pi)$ and $v(\epsilon) = 1$. We also define a valuation on N with the following standard construction : let $\mathcal{O} \subset N$ be the convex closure of \mathbb{Q} and \mathfrak{M} be the ring of infinitesimals. Then \mathcal{O} is a valuation ring, namely every element of N or its inverse lies in it. It has \mathfrak{M} as unique maximal ideal. There is therefore a valuation v' on N such that $v'(x) \geq 0$ on \mathcal{O} and $v'(x) > 0$ on \mathfrak{M} . Renaming the value group, we can set $v'(\epsilon) = 1$. Then v' extends the valuation v . As N is in the algebraic closure of $\mathbb{Q}(\epsilon, \pi)$, by standard results on valuation theory (see for example [19], Theorem 3.2.4), the value group of v' is in the divisible hull of the value group of v .

Let $b \in N$ be a positive infinitesimal. By the previous argument $v'(b)$ is rational, so there is $n \in \mathbb{N}$ such that $v'(b) > v'(\epsilon^{1/n})$. Then $v'(b/(\epsilon^{1/n})) > 0$, so $b/(\epsilon^{1/n})$ is infinitesimal and in particular $b < \epsilon^{1/n}$. \square

Let $A = \{x \in M : x < \pi\}$. So A is an externally definable initial segment of M . Consider the externally definable set $X = \{(x, y) \in M^2 : x \in A \wedge y \notin A\}$. Let $\phi(x, y; t) = (x < t \wedge y > t)$. Then $\phi(x, y; \pi)$ is a definition of X . However it is not an honest definition because it is not included in the M -definable set $\{(x, y) : y - x > \epsilon\}$. We actually show more.

Claim 1 : There is no honest definition of X with parameters in N .

Proof : Assume that $\chi(x, y)$ is such a definition. Consider $c = \inf\{y - x : y - x > 0 \wedge \chi(x, y)\}$. Then $c \in N$. For every $0 < \epsilon \in M$ infinitesimal, we have $c > \epsilon$ by the same argument as above. By the previous lemma, there is $0 < e \in \mathbb{Q}$ such that $c > e$. This is absurd as $\chi(x, y) \supseteq X$.

Let p be the global 1-type such that for $a \in \mathbb{M}$, $p \vdash x > a$ if and only if there is $b \in A \subset M$ such that $a < b$. Thus p is finitely satisfiable in M . Let $a_0 = \pi$ and $a_1 \models p|_N$. Consider the formula $\psi(x, y; a_0, a_1) = (x < a_1 \wedge y > a_0)$.

Claim 2 : The formula ψ is an honest definition of X .

Proof : Let $\theta(x, y) \in L(M)$ be a definable set. Assume that $X \subseteq \theta(M^2)$ and for a contradiction that $M \models (\exists x, y)\psi(x, y; a_0, a_1) \wedge \neg\theta(x, y)$. As p is finitely satisfiable in M , there is $u_0 \in M$ such that $\models (\exists x, y)x < u_0 \wedge y > a_0 \wedge \neg\theta(x, y)$. Consider the M -definable set $\{v : (\exists x, y)x < u_0 \wedge y > v \wedge \neg\theta(x, y)\}$. By o-minimality, this set has a supremum $m \in M \cup \{+\infty\}$. We know $m \geq a_0$, so necessarily there is $v_0 \in M$, $v_0 \notin A$ such that $M \models (\exists x, y)x < u_0 \wedge y > v_0 \wedge \neg\theta(x, y)$. This contradicts the fact that $X \subseteq \theta(M^2)$.

We therefore see that if $\phi(x, y; a)$ is a formula and M a model, then one cannot in general obtain an honest definition of $\phi(M^2; a)$ with the same parameter a . We conjecture that one can find such an honest definition with parameters in a Morley sequence of any coheir of $\text{tp}(a/M)$.

As an application, we give another proof of Shelah's expansion theorem from [56].

Proposition 2.2.9. (*T is NIP*) *Let $X \subseteq M^k$ be an externally definable set and f an M -definable function. Then $f(X)$ is externally definable.*

18 CHAPITRE 2. ENSEMBLES EXTÉRIEUREMENT DÉFINISSABLES

Démonstration. Let $\phi(x, c)$ be an honest definition of X . We show that

$$\theta(y, c) = (\exists x)(\phi(x, c) \wedge f(x) = y)$$

is a definition of $f(X)$. First, as $\phi(x, c)$ is a definition of X , we have $f(X) \subseteq \theta(M, c)$. Conversely, consider a tuple $a \in M^k \setminus f(X)$. Let $\psi(x) = (f(x) \neq a)$. Then $X \subseteq \psi(M)$. So by definition of an honest definition, $M \models \phi(x, c) \rightarrow \psi(x)$. This implies that $M \models \neg\theta(a, c)$. Thus $\theta(M, c) \subseteq f(X)$.

In fact one can check that $\theta(y, c)$ is an honest definition of $f(X)$. \square

Corollary 2.2.10 (Shelah's expansion theorem). *Let $M \models T$, be NIP and let M^{Sh} denote the expansion of M where we add a predicate for all externally definable sets of M^k , for all k . Then M^{Sh} has elimination of quantifiers in this language and is NIP.*

Démonstration. Elimination of quantifiers follows from the previous proposition, taking f to be a projection. As T is NIP, it is clear that all quantifier free formulas of M^{Sh} are dependent. It follows that M^{Sh} is dependent. \square

Note that there is an asymmetry in the notion of an honest definition. Namely if $\theta(x)$ is an honest definition of some $X \subset M$, then $\neg\theta(x)$ is not in general an honest definition of $M \setminus X$. We do not know about existence of *symmetric* honest definitions which would satisfy this. All we can do is have an honest definition contain one (or indeed finitely many) uniformly definable family of sets. This is the content of the next proposition.

Proposition 2.2.11. *(T is NIP) Let $X \subseteq M^k$ be externally definable. Let $\zeta(x, y) \in L$. Define $\Omega = \{y \in M : \zeta(M, y) \subseteq X\}$. Assume that $\bigcup_{y \in \Omega} \zeta(M, y) = X$.*

Then there is a formula $\theta(x, y)$ and $d \in M$ such that :

1. $\theta(x, d)$ is an honest definition of X ,
2. $M \models \zeta(x, c) \rightarrow \theta(x, d)$ for every $c \in \Omega$,
3. For any $c_1, \dots, c_n \in \Omega$, there is $d' \in M$ such that $\theta(M, d') \subseteq X$, and $\zeta(x, c_i) \rightarrow \theta(x, d')$ holds for all i .

Démonstration. Let $M \prec N$ where N is $|M|^+$ -saturated. Consider the set $Y \subset M$ defined by

$$y \in Y \iff (\forall x \in M)(\zeta(x, y) \rightarrow x \in X).$$

By Corollary 2.2.10, this is an externally definable subset of M , so there is $\psi(x) \in L(N)$ a definition of it. Let also $\phi(x) \in L(N)$ be a definition of X . Let $(N, M) \prec (N', M')$ be an elementary extension of the pair, sufficiently saturated. Applying Proposition 2.2.1 with $p(y) = \{\mathbf{P}(y)\}$, $A = M$ we obtain a formula $\alpha(y, d) \in L(M')$ such that $\alpha(M, d) = \psi(M)$ and $N' \models \alpha(y, d) \rightarrow^{\mathbf{P}(y)} \psi(y)$. Set $\theta(x, d) = (\exists y)(\alpha(y, d) \wedge \zeta(x, y))$. We check that $\theta(x, d)$ satisfies the required properties.

First, let $a \in M'$ such that $N' \models \theta(a, d)$. Then as $M' \prec N'$, there is $y_0 \in M'$ such that $\alpha(y_0, d) \wedge \zeta(a, y_0)$. By construction of $\alpha(y, d)$, this implies that $N' \models \psi(y_0)$. So by definition of $\psi(y)$, $N' \models \phi(a)$, so $N' \models \theta(x, d) \rightarrow^{\mathbf{P}(x)} \phi(x)$. Now, assume that $a \in X$. By

hypothesis, there is $y_0 \in \Omega$ such that $M \models \zeta(a, y_0)$. Then $\psi(y_0)$ holds, and as $y_0 \in M$, $N' \models \alpha(y_0, d)$. Therefore $N' \models \theta(a, d)$. This proves that $\theta(x, d)$ is an honest definition of X .

Next, if $c \in \Omega$, then $N' \models \alpha(c, d)$, so $N' \models \zeta(x, c) \rightarrow \theta(x, d)$.

Finally, let $c_1, \dots, c_n \in \Omega$. Then

$$N' \models (\exists d \in \mathbf{P})(\bigwedge \zeta(x, c_i) \rightarrow^{\mathbf{P}(x)} \theta(x, d)) \wedge (\theta(x, d) \rightarrow^{\mathbf{P}(x)} \phi(x)).$$

By elementarity, (N, M) also satisfies that formula. This gives us the required d' . \square

Note in particular that the hypothesis on $\zeta(x, y)$ is always satisfied for $\zeta(x, y) = (x = y)$. As an application, we obtain that large externally definable sets contain infinite definable sets.

Corollary 2.2.12. (*T is NIP*) *Let $X \subseteq M^k$ be externally definable, then if one of the two following conditions is satisfied, X contains an infinite M-definable set.*

1. *X is infinite and T eliminates the quantifier \exists^∞ .*
2. *$|X| \geq \beth_\omega$.*

Démonstration. Let $\theta(x, y)$ be the formula given by the previous proposition using $\zeta(x, y) = "x = y"$.

If the first assumption holds, then there is n such that for every $d \in M$, if $\theta(M, d)$ has size at least n , it is infinite. Take $c_1, \dots, c_n \in X$ and $d' \in M$ given by the third point of 2.2.11. Then $\theta(M, d')$ is an infinite definable set contained in X .

Now assume that $|X| \geq \beth_\omega$. By NIP, there is Δ a finite set of formulas and n such that if $(a_i)_{i < \omega}$ is a Δ -indiscernible sequence and $d \in M$, there are at most n indices i for which $\neg(\theta(a_i, d) \leftrightarrow \theta(a_{i+1}, d))$. By the Erdős-Rado theorem, there is a sequence $(a_i)_{i < \omega_1}$ in X which is Δ -indiscernible. Define $c_i = a_{\omega, i}$ for $i = 0, \dots, n$ and let d' be given by the third point of Proposition 2.2.11. Then $\theta(x, d')$ must contain an interval $\langle a_i : \omega \times k \leq i \leq \omega \times k + 1 \rangle$ for some $k \in \{0, \dots, n-1\}$. In particular it is infinite. \square

This property does not hold in general. For example in the random graph, for any κ it is easy to find a model M and $A \subset M$, $|A| \geq \kappa$ such that every M -definable subset of A is finite, while A itself is externally definable.

Also, taking $M = (\mathbb{N} + \mathbb{Z}, <)$ and $X = \mathbb{N}$ shows that $|X|$ has to be bigger than \aleph_0 in 2.2.12 in general.

Question 2.2.13. *Is it possible to replace \beth_ω by \aleph_1 in 2.2.12?*

§2.3 On dependent pairs

Setting

In this section, we assume that T is NIP. We consider a pair (M, A) with $M \models T$. If $\phi(x, a)$ is some formula of $L_P(M)$, then an *honest definition of $\phi(x, a)$ over A* is a formula $\theta(x, c) \in L_P$, $c \in P(\mathbb{M})$ such that $\theta(A, c) = \phi(A, a)$ and $\models (\forall x \in P)(\theta(x, c) \rightarrow \phi(x, a))$.

(Note that if $M \models T$, $\phi(x, c) \in L(\mathbb{M})$ and $X = \phi(M, c)$, then an honest definition of $\phi(x, c)$ over M in the pair (\mathbb{M}, M) which happens to be an L -formula is an honest definition of X in the sense of Definition 2.2.6.)

We say that an L_P -formula is *bounded* if it is of the form $Q_1 y_1 \in P \dots Q_n y_n \in P \phi(x, y_1, \dots, y_n)$ where $Q_i \in \{\exists, \forall\}$ and $\phi(x, \bar{y})$ is an L -formula, and let L_P^{bdd} be the collection of all bounded formulas. We say that T_P is bounded if every formula is equivalent to a bounded one.

Recall that a formula $\phi(x, y) \in L_P$ is said to be *NIP over $P(x)$* if there is no L_P -indiscernible (equivalently L -indiscernible if $\phi \in L$) sequence $(a_i)_{i < \omega}$ of points of P and y such that $\phi(a_i, y) \Leftrightarrow i$ is even. If this is the case, then Proposition 2.2.1 applies and in particular there is an honest definition of $\phi(x, a)$ over P for all a .

We say that T (or T_P) is *NIP over P* if every L (resp. L_P) formula is.

Given a small subset of the monster A and a set of formulas Ω (possibly with parameters) we let $A_{\text{ind}(\Omega)}$ be the structure with domain A and a relation added for every set of the form $A^n \cap \phi(\bar{x})$, where $\phi(\bar{x}) \in \Omega$.

Notice that $A_{\text{ind}(L_P^{\text{bdd}})}$ eliminates quantifiers, while $A_{\text{ind}(L)}$ not necessarily does. However $A_{\text{ind}(L_P^{\text{bdd}})}$ and $A_{\text{ind}(L)}$ are bi-interpretable.

Lemma 2.3.1. *Assume that $\varphi(xy, c) \in L_P$ has an honest definition $\vartheta(xy, d) \in L_P$ over A . Then $\theta(x, d) = (\exists y \in P)\vartheta(xy, d)$ is an honest definition of $\phi(x, c) = (\exists y \in P)\varphi(xy, c)$ over A .*

Démonstration. For $a \in P$, $\theta(a, d) \Rightarrow \vartheta(ab, d)$ for some $b \in P \Rightarrow \varphi(ab, c)$ (as $\vartheta(xy, d)$ is honest and $ab \in P \Rightarrow \phi(a, c)$).

For $a \in A$, $\phi(a, c) \Rightarrow \varphi(ab, c)$ for some $b \in A \Rightarrow \vartheta(ab, d)$ (as $\vartheta(A, d) = \phi(A, c)$) $\Rightarrow \theta(a, d)$. \square

We will be using λ -big models (see [28, 10.1]). We will only use that if N is λ -big, then it is λ -saturated and strongly λ -homogeneous (that is, for every $\bar{a}, \bar{b} \in N^{<\lambda}$ such that $(N, \bar{a}) \equiv (N, \bar{b})$ there is an automorphism of N taking \bar{a} to \bar{b}) (see [28, 10.1.2 + Exercise 10.1.4]). Every model M has a λ -big elementary extension N .

Lemma 2.3.2. *1) If $N \succeq M$, M is ω -big, N is $|M|^+$ -big, and $a, b \in M^{<\omega}$ then $\text{tp}_L(a) = \text{tp}_L(b) \Leftrightarrow \text{tp}_{L_P}(a) = \text{tp}_{L_P}(b)$ in the sense of the pair (N, M) .*

2) Let $\phi(x, y) \in L_P$, (M, A) ω -big, $(a_i)_{i < \omega} \in M^\omega$ be L_P -indiscernible, and let $\theta(x, d_0)$ be an honest definition for $\phi(x, a_0)$ over A (where d_0 is in P of the monster model). Then we can find an L_P -indiscernible sequence $(d_i)_{i < \omega} \in P^\omega$ such that $\theta(x, d_i)$ is an honest definition for $\phi(x, a_i)$ over A .

Démonstration. 1) We consider here the pair (N, M) as an L_P -structure, where $P(x)$ is a new predicate interpreted in the usual way. Let $\sigma \in \text{Aut}_L(M)$ be such that $\sigma(a) = b$. As N is big, it extends to $\sigma' \in \text{Aut}_L(N)$, with $\sigma'(M) = M$. But then actually $\sigma' \in \text{Aut}_{L_P}(N)$ (since it preserves all L -formulas and P).

2) Let $(N, B) \succeq (M, A)$ be $|M|^{+}$ -big. We consider the pair of pairs $\text{Th}((N, B), (M, A))$ in the language $L_{P, P'}$, with $P'(N) = M$. By 1) the sequence $(a_i)_{i < \omega}$ is $L_{P, P'}$ -indiscernible. The fact that $\theta(x, d_0)$ is an honest definition of $\phi(x, a_0)$ over A is expressible by the formula

$$(d_0 \in P) \wedge ((\forall x \in P' \cap P) \theta(x, d_0) \equiv \phi(x, a_0)) \wedge ((\forall x \in P) \theta(x, d_0) \rightarrow \phi(x, a_0)).$$

By $L_{P, P'}$ -indiscernibility, for each i , we can find d_i such that the same formula holds of (a_i, d_i) . Then using Ramsey, for any finite $\Delta \subset L_P$, we can find an infinite subsequence $(a_i, d_i)_{i \in I}$, $I \subseteq \omega$ that is Δ -indiscernible. As (a_i) is indiscernible, we can assume $I = \omega$. Then by compactness, we can find the d_i 's as required. \square

We will need the following technical lemma.

Lemma 2.3.3. *Let $(M, A) \models T_P$ be ω -big and assume that $A_{\text{ind}(L_P)}$ is NIP.*

Let $(a_i)_{i < \omega} \in M^\omega$ be L_P -indiscernible, $(b_{2i})_{i < \omega} \in A^\omega$ and $\Delta((x_i)_{i < n}; (y_i)_{i < n}) \in L_P$ be such that $\Delta((x_i)_{i < n}; (a_i)_{i < n})$ has an honest definition over A by an L_P -formula, and $\models \Delta(b_{2i_0}, \dots, b_{2i_{n-1}}; a_{2i_0}, \dots, a_{2i_{n-1}})$ for any $i_0, \dots, i_{n-1} < \omega$.

Then there are $i_0, \dots, i_{n-1} \in \omega$ with $i_j \equiv j \pmod{2}$ and $(b_{i_j})_{j \equiv 1 \pmod{2}, < n} \in P$ such that $\models \Delta(b_{i_0}, \dots, b_{i_{n-1}}; a_{i_0}, \dots, a_{i_{n-1}})$.

Démonstration. To simplify notation assume that n is even. Let

$$\Delta'((x_{2i})_{2i < n}; (y_i)_{i < n}) = (\exists x_1 x_3 \dots x_{n-1} \in P) \Delta((x_i)_{i < n}; (y_i)_{i < n}).$$

By assumption and Lemma 2.3.1 $\Delta'((x_{2i})_{2i < n}; (a_i)_{i < n})$ has an honest definition over A by some L_P -formula, say $\theta((x_{2i})_{2i < n}, d)$ with $d \in P$. Since $A_{\text{ind}(L_P)}$ is NIP, let $N = \text{alt}(\theta)$ inside P .

Choose even $i_0, i_2, \dots, i_{n-2} \in \omega$ such that $i_{j+2} - i_j > N$ and consider the sequence $(\bar{a}_i)_{0 < i < N}$ with $\bar{a}_i = a_{i_0} a_{i_0+i} a_{i_2} a_{i_2+i} \dots a_{i_{n-2}} a_{i_{n-2}+i}$. It is L_P -indiscernible (and extends to an infinite L_P -indiscernible sequence). By Lemma 2.3.2 we can find an L_P -indiscernible sequence $(d_i)_{i < N}$, $d_i \in P$ such that $\theta((x_{2i})_{2i < n}; d_i)$ is an honest definition for $\Delta'((x_{2i})_{2i < n}; \bar{a}_i)$. By assumption $\theta((b_{i_{2j}})_{2j < n}; d_i)$ holds for all even $i < N$. But then since $N = \text{alt}(\theta)$ inside P , it must hold for some odd $i' < N$. By honesty this implies that $\Delta'((b_{i_{2j}})_{2j < n}; \bar{a}_{i'})$ holds, and decoding we find some $(b_{i_{2j}+i'})_{2j < n} \in P^{\frac{n}{2}}$ as wanted. \square

Now the main results of this section.

Theorem 2.3.4. *Assume T is NIP and T_P is NIP over P . Then every bounded formula is NIP.*

Démonstration. We prove this by induction on adding an existential bounded quantifier (since NIP formulas are preserved by boolean operations). So assume that $\phi(x, y) = (\exists z \in \mathbf{P}) \psi(xz, y)$ has IP, where $\psi(xz, y) \in L_{\mathbf{P}}^{\text{bdd}}$ is NIP. Then there is an ω -big $(M, A) \models T_{\mathbf{P}}$ and an $L_{\mathbf{P}}$ -indiscernible sequence $(a_i)_{i < \omega} \in M^\omega$ and $c \in M$ such that $\phi(a_i, c) \Leftrightarrow i \equiv 0 \pmod{2}$. Then we can assume that there are $b_{2i} \in A$ such that $(a_{2i}b_{2i})$ is $L_{\mathbf{P}}$ -indiscernible and $\models \psi(a_{2i}b_{2i}, c)$.

Notice that from $T_{\mathbf{P}}$ being NIP over \mathbf{P} it follows that $A_{\text{ind}(L_{\mathbf{P}})}$ is NIP and that every $L_{\mathbf{P}}$ -formula has an honest definition over A . For $\delta \in L_{\mathbf{P}}$ take $\Delta_\delta((x_i)_{i < n}; (y_i)_{i < n})$ to be an $L_{\mathbf{P}}$ -formula saying that $(x_i y_i)_{i < n}$ is δ -indiscernible. Applying Lemma 2.3.3, we obtain $i_0, \dots, i_n \in \omega$ with $i_j \equiv j \pmod{2}$ and $(b_{i_j})_{j \equiv 1 \pmod{2}, < n} \in \mathbf{P}$ such that $(a_{i_k} b_{i_k})_{k < n}$ is δ -indiscernible. Since $\models \neg(\exists z \in \mathbf{P}) \psi(a_{2i+1}z, c)$ for all i , we see that $\psi(a_{i_k} b_{i_k}, c)$ holds if and only if k is even. Taking n and δ large enough, this contradicts dependence of $\psi(xz, y)$. \square

Corollary 2.3.5. *Assume T is NIP, $A_{\text{ind}(L)}$ is NIP and $T_{\mathbf{P}}$ is bounded. Then $T_{\mathbf{P}}$ is NIP.*

Démonstration. Since $A_{\text{ind}(L_{\mathbf{P}}^{\text{bdd}})}$ is interpretable in $A_{\text{ind}(L)}$ the hypothesis implies that $A_{\text{ind}(L_{\mathbf{P}}^{\text{bdd}})}$ is NIP. Thus, if $\bar{a} = (a_i)_{i < n}$ is a sequence inside \mathbf{P} then any $\Delta(\bar{x}, \bar{a})$ has an honest definition over A (although we don't yet know that $\Delta(\bar{x}, \bar{y})$ is NIP over \mathbf{P} , we do know that $\Delta(\bar{x}, \bar{a})$ is NIP over \mathbf{P} , so Proposition 2.2.1 applies). We can then use the same proof as in 2.3.4 to ensure that $T_{\mathbf{P}}$ is NIP over \mathbf{P} , and finally apply Theorem 2.3.4 to conclude. \square

Corollary 2.3.6. *Assume T is NIP, and let (M, N) be a pair of models of T ($N \prec M$). Assume that $T_{\mathbf{P}}$ is bounded, then $T_{\mathbf{P}}$ is NIP.*

Démonstration. $N_{\text{ind}(L)}$ is dependent, and so the hypotheses of Corollary 2.3.5 are satisfied. \square

Note that the boundedness assumption cannot be dropped, because for example a pair of real closed fields can have IP, and also there is a stable theory such that some pair of its models has IP ([49]).

§2.4 Applications

In this section we give some applications of the criteria for the dependence of the pair.

2.4.1 Naming an indiscernible sequence

In [5] Baldwin and Benedikt prove the following.

Fact 2.4.1. (*T is NIP*) *Let $I \subset M$ be an indiscernible sequence indexed by a dense complete linear order, small in M (that is every $p \in S_{<\omega}(I)$ is realised in M). Then*

- 1) *$\text{Th}(M, I)$ is bounded ([5, Theorem 3.3]),*

- 2) $(M, I) \equiv (N, J)$ if and only if $EM(I) = EM(J)$ ([5, Theorem 8.1]),
 3) The L_P -induced structure on \mathbf{P} is just the equality (if I is totally transcendental) or the linear order otherwise ([5, Corollary 3.6]).

It is not stated in the paper in exactly this form because the bounded formula from [5, Theorem 3.3] involves the order on the indiscernible sequence. However, it is not a problem. If the sequence $I = (a_i)$ is not totally indiscernible, then the order is L -definable (maybe after naming finitely many constants). Namely, we will have $\phi(a_0, \dots, a_k, a_{k+1}, \dots, a_n) \wedge \neg \phi(a_0, \dots, a_{k+1}, a_k, \dots, a_n)$ for some $k < n$ and $\phi \in L$ (as the permutation group is generated by transpositions). But then the order on I is given by $y_1 < y_2 \leftrightarrow \phi(a'_0 \dots a'_{k-1}, y_1, y_2, a'_{k+2}, \dots, a'_n)$, for any $a'_0 \dots a'_{k-1} I a'_{k+2} \dots a'_n$ indiscernible (and we can find such $a'_0 \dots a'_{k-1} I a'_{k+2} \dots a'_n$ in M by the smallness assumption). If I is an indiscernible set, then the stable counterpart of their theorem [5, 3.3] applies giving a bounded formula using just the equality (as the proof in [5, Section 4] only uses that for an NIP formula $\phi(x, y)$ and an arbitrary c , $\{a_i : \phi(a_i, c)\}$ is either finite or cofinite, with size bounded by $\text{alt}(\phi)$).

The following answers Conjecture 9.1 from that paper.

Proposition 2.4.2. *Let (M, I) be a pair as described above, obtained by naming a small, dense, complete indiscernible sequence. Then T_P is NIP.*

Démonstration. By 1) and 3) above, all the assumptions of Corollary 2.3.5 are satisfied. \square

It also follows that every unstable dependent theory has a dependent expansion with a definable linear order.

Recall the following definition (one of the many equivalent) from [53].

Definition 2.4.3. [53, Observations 2.1 and 2.10] T is strongly (resp. strongly⁺) dependent if for any infinite indiscernible sequence $(\bar{a}_i)_{i \in I}$ with $\bar{a}_i \in M^\omega$, I a complete linear order, and finite tuple c there is a finite $u \subset I$ such that for any two $i_1 < i_2 \in u$, $(i_1, i_2) \cap u = \emptyset$ the sequence $(\bar{a}_i)_{i \in (i_1, i_2)}$ is indiscernible over c (resp. $c \cup (\bar{a}_i)_{i \in (-\infty, i_1] \cup [i_2, \infty)}$).

T is dp -minimal (resp. dp^+ -minimal) when for a singleton c there is such a u of size 1.

For a general NIP theory, the property described in the definition holds, but with $u \subset I$ of size $|T|$, instead of finite. We can take u to be the set of *critical points* of I defined by : $i \in I$ is critical for a formula $\phi(x; y_1, \dots, y_n, c) \in L$ if there are $j_1, \dots, j_n \neq i$ such that $\phi(a_i; a_{j_1}, \dots, a_{j_n}, c)$ holds, but in every open interval of I containing i , we can find some i' such that $\neg \phi(a_{i'}; a_{j_1}, \dots, a_{j_n}, c)$ holds. One can show (see [2, Section 3]) that given such a formula $\phi(x; y_1, \dots, y_n, c)$, the set of critical points for ϕ is finite. Also T is strongly⁺ dependent if and only if for every finite set c of parameters, the total number of critical points for formulas in $L(c)$ is finite.

Unsurprisingly dp -minimality is not preserved in general after naming an indiscernible sequence. By [21, Lemma 3.3] in an ordered dp -minimal group, there is no infinite definable nowhere-dense subset, but of course every small indiscernible sequence is like this.

There are strongly dependent theories which are not strongly⁺ dependent, for example p -adics ([53]). In such a theory, strong dependence is not preserved by naming an indiscernible sequence.

Proposition 2.4.4. *Let T be not strongly⁺ dependent, witnessed by a dense complete indiscernible sequence $(\bar{a}_i)_{i \in I}$ of finite tuples. Let \mathbf{P} name that sequence in a big saturated model. Then $T_{\mathbf{P}}$ is not strongly dependent.*

Démonstration. So let $(\bar{a}_i)_{i \in I}, c$ witness failure of strong⁺ dependence. By dependence of T , let $u \subset I$ be chosen as above. Notice that for every $\phi(x; y_1, \dots, y_n, c)$, the finite set of its critical points in I is $L_{\mathbf{P}}$ -definable over c (and possibly finitely many parameters, using order on I in the non-totally indiscernible case, and just the equality otherwise). As in our situation u is infinite, we get infinitely many different finite subsets of $(\bar{a}_i)_{i \in I}$ definable over c , in $T_{\mathbf{P}}$. As $(\bar{a}_i)_{i \in I}$ is still indiscernible in $T_{\mathbf{P}}$ by Fact 2.4.1, 3), this contradicts strong dependence. \square

Question 2.4.5. *Is strong⁺ dependence preserved by naming an indiscernible sequence?*

2.4.2 Dense pairs and related structures

Van den Dries proves in [66] that in a dense pair of o-minimal structures, formulas are bounded. This is generalised in [8] to lovely pairs of geometric theories of thorn-rank 1. From Theorem 2.3.6, we conclude that such pairs are dependent.

This was already proved by Berenstein, Dolich and Onshuus in [9] and generalised by Boxall in [11]. Our result generalises [9, Theorem 2.7], since the hypothesis there (acl is a pregeometry and A is “innocuous”) imply boundedness of $T_{\mathbf{P}}$. To see this take any two tuples a and b and assume that they have the same bounded types. Let $a' \in \mathbf{P}$ be such that aa' is a \mathbf{P} -independent tuple. Then by hypothesis, we can find b' such that $\text{tp}_{L_{\mathbf{P}}^{\text{bdd}}}(\bar{b}\bar{b}') = \text{tp}_{L_{\mathbf{P}}^{\text{bdd}}}(\bar{a}\bar{a}')$. Now the fact that aa' is \mathbf{P} -independent can be expressed by bounded formulas. In particular $\bar{b}\bar{b}'$ is also \mathbf{P} -independent. So by innocuous, $\text{tp}_{L_{\mathbf{P}}}(\bar{a}\bar{a}') = \text{tp}_{L_{\mathbf{P}}}(\bar{b}\bar{b}')$ and we are done.

It is not clear to us if Boxall’s hypothesis imply that formulas are bounded. (However, note that in the same paper Boxall applies his theorem to the structure of \mathbb{R} with a named subgroup studied by Belegardek and Zilber, where we know that formulas are bounded.)

The paper [9] gives other examples of theories of pairs for which formulas are bounded, including dense pairs of p -adic fields and weakly o-minimal theories, recast in the more general setting of *geometric topological structures*.

Similar theorems are proved by Günaydin and Hieronymi in [25]. Their Theorem 1.3 assumes that formulas are bounded along with other hypothesis, so is included in Theorem 2.3.6. They apply it to show that pairs of the form (\mathbb{R}, Γ) are dependent, where

$\Gamma \subset \mathbb{R}^{>0}$ is a dense subgroup with the *Mann property*. We refer the reader to [25] for more details.

In this same paper the authors also consider the case of tame pairs of o-minimal structures. This notion is defined and studied in [67]. Let T be an o-minimal theory. A pair (N, M) of models of T is *tame* if $M \prec N$ and for every $a \in N$ which is in the convex hull of M , there is $st(a) \in M$ such that $|a - st(a)| < b$ for every $b \in M^{>0}$. It is proved in [67] that formulas are bounded in such a pair, so again it follows from Theorem 2.3.6 that T_P is dependent. Note that Günaydin and Hieronymi prove this using their Theorem 1.4 involving quantifier elimination in a language with a new function symbol. This theorem does not seem to factorise trivially through 2.3.5. They also prove in that same paper that the pair $(\mathbb{R}, 2^{\mathbb{Z}})$ is dependent.

Let C be an elliptic curve over the reals, defined by $y^2 = x^3 + ax + b$ with $a, b \in \mathbb{Q}$, and let $P \subseteq \mathbb{Q}^2$ name the set of its rational points. This theory is studied in [24], where it is proved in particular that

Fact 2.4.6. 1) $\text{Th}(\mathbb{R}, C(\mathbb{Q}))$ is bounded (follows from [24, Theorem 1.1])
 2) $A_{\text{ind}(L_P)}$ is NIP (follows from [24, Proposition 3.10])

Applying Corollary 2.3.5 we conclude that the pair is dependent.

Mesures génériquement stables

écrit en commun avec Ehud Hrushovski et Anand Pillay

§3.1 Introduction and preliminaries

In this introduction we will discuss background and motivation, describe and summarize our main results, and then recall some essential definitions and prior results. A familiarity with the earlier papers [32] and [33] would be advantageous, but we will try to make the bulk of the paper accessible to a wider audience, even though we are somewhat terse. Even in the introduction we may make some rather advanced comments or references, and the general reader should feel free to ignore these at least on the first reading.

Shelah defined a formula $\phi(x, y)$ to have the *independence property* if there exist arbitrarily large (finite) sets A such that *any* subset B of A has the form $\{a \in A : \phi(a, b)\}$, for some parameter b . A theory has NIP if no formula has the independence property. An equivalent definition in a combinatorial / probabilistic rather than logical setting was found by Chervonenkis and Vapnik [68]. α -minimal and p -minimal theories are notable examples.

A general theme in this paper is “stable-like” behaviour in theories with NIP. One of the main points is to develop the theory of “generically stable measures” in NIP theories, in analogy with generically stable types. A “generically stable type” is a global type (namely a complete type over a saturated model) which looks very much like a type in a stable theory, for example it is both definable over and finitely satisfiable in some small model M . The theory, at least in the NIP context was developed in [56], [33] and [65]. Among the consequences (or even equivalences) of generic stability of a type p , assuming T has NIP, are *nonforking symmetry* (or the total indiscernibility of any “Morley sequence” in p), as well as *stationarity*, in the sense that p is the *unique* global nonforking extension of its restriction to M .

In the theory of algebraically closed valued fields generically stable types coincide with stably dominated types and play a major role in the structural analysis of definable sets [27] as well as in a model-theoretic approach to Berkovich spaces [31]. However in \mathbf{o} -minimal theories and \mathbf{p} -adically closed fields for example, there are *no* (nonalgebraic) generically stable types.

On the other hand, what we have called Keisler measures (introduced in Keisler's seminal paper [37]), are the natural generalization of complete types to finitely additive $[0, 1]$ valued measures on Boolean algebras of definable sets. Keisler showed (in slightly different terms) that in a NIP theory, for any Keisler measure μ on a model M any formula $\phi(x, y)$ and any $\epsilon > 0$, there exist finitely many formulas $\phi(x, b_i)$ such that for any b , $\mu(\phi(x, b_i) \triangle \phi(x, b)) < \epsilon$ for some i . To see this, take a maximal set $\{b_i\}$ such that $\mu(\phi(x, b_i) \triangle \phi(x, b_j)) \geq \epsilon/2$ for $i \neq j$. If this set is finite, we are done. If it is infinite, by compactness one obtains an indiscernible sequence $(b_n : n \in \mathbb{N})$ and some measure μ' with the same property. So $\mu(\phi(x, b_m) \setminus \phi(x, b_{m+1})) \geq \epsilon/4$ for all odd m (or for all even m ; say odd.) It follows by elementary measure considerations that $(\mu(\phi(x, b_m) \setminus \phi(x, b_{m+1})) : m = 2, 4, \dots)$ cannot be k -inconsistent, for any k . So $\{\phi(x, b_m) : m = 1, 2, \dots\} \cup \{\neg\phi(x, b_m) : m = 2, 4, \dots\}$ is consistent. But by indiscernibility the same must be true for any subset in place of the odds, contradicting NIP.

Keisler measures play an important role in the solution of certain conjectures on groups in \mathbf{o} -minimal structures [32]. They were studied further and from a more stability-theoretic point of view in [33]. In fact in the latter paper, we *defined* generically stable measures to be global Keisler measures which are both definable over and finitely satisfiable in some small model. We also found natural examples as translation invariant measures on suitable definable groups (such as definably compact groups in \mathbf{o} -minimal theories). However, there were on the face of it technical obstacles to obtaining analogous properties (like stationarity, total indiscernibility) for generically stable measures as for generically stable types. For example, what is a “realization” of a measure, or a “Morley sequence in a measure”? This is solved in various ways in the current paper, including making heavy use of Keisler's “smooth measures” (see section 2). Essentially a complete counterpart to the type case is obtained, the main results along these lines being Theorem 3.2 and Proposition 3.3, where another property, “frequency interpretation measure” makes an appearance. Moreover we also point out how widespread generically stable measures are in NIP theories.

Let us take the opportunity to remark that a natural formal way to deal with “technical” issues such as realizing Keisler measures would be to pass to the randomization T^R of T . T^R is a continuous first order theory whose models are random variables in models of T . The type spaces of T^R correspond to the spaces of Keisler measures (over \emptyset) of T . This randomization was introduced by Keisler and situated in the context of continuous logic by Ben Yaacov and Keisler [70]. Ben Yaacov proved that T^R has NIP if T does, and further showed that making systematic use of T^R would provide, in principle, another route to the results of the current paper ([71], [72]). Measures in NIP theories are roughly of the same complexity as types, as is evidenced for instance by boundedness of the number of formulas modulo measure zero. But measures on the space of mea-

asures appear to be genuinely analytic objects, and required nontrivial analytic tools in Ben Yaacov's treatment. We make use of a weak version of Ben Yaacov's preservation theorem (see Lemma 2.10) to give one proof of our characterization of generically stable measures (Theorem 3.2), but also give an independent proof remaining within the usual model theoretic framework.

In section 4 we generalize the notion of a *group with finitely satisfiable generics* or with the *fsg* property, to types and measures, and make the connection with generic stability.

In section 5 we introduce a weak notion of “compact domination” where the set being dominated is a space of types rather than a definable or type-definable set. We relate this to stationarity of measures (unique nonforking extensions) in what we consider to be a measure-theoretic version of the finite equivalence theorem.

In section 6, we prove smoothness (unique extension to a global Keisler measure) of Borel probability measures on real or p -adic semialgebraic sets, yielding a quite extensive strengthening of work by Karpinski and Macintyre in the case of Haar measure.

As far as sections 2, 3 and 6 are concerned, the paper is relatively self-contained. However sections 4 and 5 make rather more references to the earlier papers [32] and [33], and not only hyperimaginaries but also the compact Lascar group are involved.

The current paper does not only follow on from those two earlier papers, but also naturally continues and builds on Keisler's original papers [37], [36].

We fix a complete first order theory T . We typically work in T^{eq} . For convenience we choose a very saturated “monster model” or “universal domain” $\bar{M} = \bar{M}^{eq}$. M, N, M_0, \dots denote small elementary submodels. For now A, B, C, \dots denote subsets, usually small, of \bar{M} . x, y, \dots range over (finitary) variables and by convention a variable carries along with it its sort.

The reader is referred to say [47], [2], [50], [54] as well as [32], [33], for extensive and detailed material around stable theories, NIP theories as well as the adaptation/interpretation of forking to types and measures in NIP theories.

However we recall here the key notions relevant to the current paper.

It is convenient to start with the notion of a finitely additive measure μ on an arbitrary Boolean algebra Ω : $\mu(b) \in [0, 1]$ for all b in Ω , $\mu(1) = 1, \mu(0) = 0$ and μ is finitely additive. As in section 4 of [33], such a measure on a Boolean algebra Ω can be identified with a regular Borel probability measure on the Stone space S_Ω of Ω . The set of finitely additive measures on Ω is naturally a compact space.

We apply this to our monster model \bar{M} . By a Keisler measure μ_x over A we mean a finitely additive measure on the Boolean algebra of formulas $\phi(x)$ over A up to equivalence in \bar{M} . So a Keisler measure over A generalizes the notion of a complete type over A rather than a partial type over A . By a global Keisler measure we mean one over \bar{M} . So again a global Keisler measure generalizes the notion of a global complete type. We repeat from the previous paragraph that a Keisler measure μ_x over A coincides with a regular probability measure on $S_x(A)$. We often talk about closed, open, Borel, sets, over A . So for example, a Borel set over A is simply the union of the sets of realizations in \bar{M} of types $p \in S(A)$, for p in some given Borel subset of $S_x(A)$.

Definition 3.1.1. Let μ_x be a Keisler measure over A , and $p \in S(A)$, then we say that p is *weakly random* for μ if $\mu(\phi(x)) > 0$ whenever $p \vdash \phi(x)$ (where ϕ is a formula over A).

Similarly, a is weakly random for μ (over A) if $\text{tp}(a/A)$ is weakly random.

Keisler [37] uses the notion of a measure over or on a *fragment*, which it is now convenient to work with in a generalized form. By a fragment F he means a small collection of formulas $\phi(x)$ (or definable sets of sort x) which is closed under (finite) Boolean combination. (A typical case is the collection of all formulas over a given base set A .) Then a Keisler measure on or over F is simply a finitely additive probability measure on this Boolean algebra of definable subsets of sort x . As above this identifies with a regular Borel probability measure on the space S_F of complete types over F . He also remarks that if $F \subseteq G$ are fragments (in sort x) then any Keisler measure on F extends to one on G . In particular any Keisler measure on F extends to a global Keisler measure on the sort of x .

For most of this paper this notion of fragment is adequate, and the reader may proceed with this in mind, at least until section 5. However in some situations we will need to consider algebras of subsets of \bar{M} that, while contained in the Borel subalgebra of S_F for various fragments F of formulas, cannot canonically be presented in this manner. We therefore give in advance a formalism beginning with closed rather than clopen sets, i.e. partial types rather than formulas. Our fragments correspond to small topological quotients of the space of global types : an element of the fragment is the pullback of a closed set. We describe this more syntactically in the next paragraph.

Let F now consist of a small collection of partial types $\Sigma(x)$ in a fixed set of variables x , identified if you wish with their sets of realizations in \bar{M} . We assume F is closed under finite disjunctions and (possibly infinite) conjunctions. We will call a subset of the x -sort of \bar{M} *closed over F* if it is defined by a partial type in F , and *open over F* if it is the complement of a closed over F set (and also we can obtain the Borel over F sets).

Definition 3.1.2. (a) Let F be as in the above paragraph. We call F a *fragment* if
 (i) any open set over F is a union of closed sets over F , and
 (ii) any two disjoint closed over F sets are separated by two disjoint open over F sets.
 (b) If F is a fragment, let S_F denote the set of maximal partial types in F (i.e. maximal among partial types in F).

Clearly a fragment in the sense of Keisler extends uniquely to a fragment in the sense of Definition 1.1.

For a fragment F define a topology on S_F in the obvious way : a closed set is by definition a set of points extending a given partial type in F . Then with this definition it is clear that S_F will be a compact Hausdorff space.

Definition 3.1.3. By a Keisler measure on or over a fragment F we mean a map from the set of closed/open over F sets to $[0, 1]$ which is induced by a regular Borel probability measure on the space S_F .

For hyperimaginaries, as well as the notion $\text{bdd}(A)$ (set of hyperimaginaries in the bounded closure of A) see [26] or [69].

Lemma 3.1.4. (i) *Let A be a small set of hyperimaginaries. Then the collection of partial types over A is a fragment.*
(ii) *Let $F \subseteq G$ be fragments (in sort x). Then any Keisler measure over F extends to a Keisler measure over G .*

One more definition at the level of fragments is :

Definition 3.1.5. Let μ be a measure over a fragment F . Let D be a Borel set over F with positive μ measure. Then the localization μ_D of μ at D is defined by : For any Borel E over F , $\mu_D(E) = \mu(E \cap D)/\mu(D)$.

Now we pass to forking for measures in NIP theories. First, T is said to have the independence property, if there is an indiscernible (over \emptyset) sequence $(a_i : i < \omega)$ and formula $\phi(x, b)$ such that $\models \phi(a_i, b)$ for i even, and $\models \neg \phi(a_i, b)$ for i odd. We usually say that T is (or has) NIP if T does not have the independence property.

We recall that a formula $\phi(x, b)$ (where we exhibit the parameters) *divides over* a small set A if there is an A -indiscernible sequence $(b_i : i < \omega)$ with $b_0 = b$ such that $\{\phi(x, b_i) : i < \omega\}$ is inconsistent. A formula *forks over* A if it implies a finite disjunction of formulas each of which divides over A . We say that a global Keisler measure μ_x does not divide (does not fork) over a small set A if every formula $\phi(x)$ with positive μ -measure does not divide (does not fork) over A . In fact for such global μ , not dividing over A and not forking over A are equivalent, and A can even be a set of hyperimaginaries. Recall from [33] that assuming T has NIP, μ does not fork over A iff μ is $\text{Aut}(\bar{M}/\text{bdd}(A))$ invariant (we just say $\text{bdd}(A)$ -invariant) iff μ is Borel definable over $\text{bdd}(A)$. Here Borel-definability of μ over A , means that for a given formula $\phi(x, y) \in L$ and closed subset C of $[0, 1]$, $\{b : \mu(\phi(x, b) \in C)\}$ is Borel over A . We persist in calling a global measure μ_x *definable* over A if for $\phi(x, y) \in L$ and closed $C \subseteq [0, 1]$, $\{b : \phi(x, b) \in C\}$ is closed over A , namely type-definable over A . (Although the expression ∞ -definable might be better.) We also say that μ is finitely satisfiable in A (where usually A is a model M) if every formula over \bar{M} with positive μ -measure is satisfied by some element (or tuple) from A . These are all natural generalizations of the corresponding classical notions for global types.

Let us make the important remark that if the global Keisler measure μ_x is finitely satisfiable over A , then it is also A -invariant, hence (assuming that T has NIP) is Borel definable (over A).

The (nonforking) *product* of measures μ_x and λ_y is a fundamental notion in this paper (as well as in [33]). Identifying a global Keisler measure μ_x with a measure on $S_x(\bar{M})$, then this could not simply be the usual product measure because the type space $S_{xy}(\bar{M})$ is not the product of $S_x(\bar{M})$ with $S_y(\bar{M})$ (and the same issue arises for types). In the case of types, if $p(x)$, $q(y)$ are global complete types, and $p(x)$ does not split over A for some small A (equivalently is A -invariant), then we can form $p(x) \otimes$

$q(y)$ in variables xy , defined as $\text{tp}(a, b/\bar{M})$ where b realizes q and a realizes $p|\bar{M}, b$. Equivalently, $\phi(x, y, m) \in p(x) \otimes q(y)$ if for some (any) b realizing $q|A, m$, $\phi(x, b) \in p$. Now if $\mu(x)$ is Borel definable (over A) say, and $\lambda(y)$ arbitrary (both global say) then the analogous product $\mu_x \otimes \lambda_y$ is obtained via *integration*: Pick a formula $\phi(x, y)$ over \bar{M} . For any $q(y) \in S_y(\bar{M})$, and realization b of q , we can consider the extension $\mu' = \mu|(\bar{M}, b)$ of μ' given by applying the same Borel definition. In any case $\mu'(\phi(x, b))$ depends only on q , so we can write it as $f(q)$ for some function $f : S_y(\bar{M}) \rightarrow [0, 1]$. The Borel definability of μ says that the function f is Borel (preimage of a closed set is Borel). Hence we can integrate f along λ (treated as a Borel measure on $S_y(\bar{M})$), to obtain $\int_{S_y(\bar{M})} f(q) d\lambda$. And we call this $(\mu(x) \otimes \lambda(y))(\phi(x, y))$.

Note that this integral can be “computed” as follows : again choose a formula $\phi(x, y, m)$ where now we exhibit additional parameters from \bar{M} as m . Fix natural number N and partition $[0, 1]$ into equal intervals I_1, \dots, I_N of length $1/N$, let $Y_j = \{b : \mu(\phi(x, b, m)) \in I_j\}$ (a Borel set over A, m), let c_j be the midpoint of I_j . Let $F_N = \sum_{j=1, \dots, N} \lambda(Y_j) c_j$. Then $(\mu(x) \otimes \lambda(y))(\phi(x, y, m)) = \lim_{N \rightarrow \infty} F_N$.

We will often use this, when doing approximations or computations.

We emphasize that the product $\mu_x \otimes \lambda_y$ is well defined only when μ is Borel-definable. This implies that μ does not fork over some small set, and the converse is true assuming NIP.

Lemma 3.1.6. *Suppose that $\mu(x)$, $\lambda(y)$ are global Keisler measures which are both definable. Then so is $\mu(x) \otimes \lambda(y)$. Likewise for Borel definable, and (assuming NIP) “finitely satisfiable in a small model”.*

Proof. Let us just deal with the finitely satisfiable case, the proof of which will be an elementary example of methods which pervade the paper. Assume that both μ and λ are finitely satisfiable in M . They are thus M -invariant and by NIP Borel-definable over M . It therefore makes sense to consider the product $\mu_x \otimes \lambda_y$. We show that it is itself finitely satisfiable in M . Let $\phi(x, y, m)$ be a formula over \bar{M} with positive $\mu_x \otimes \lambda_y$ measure (where we exhibit the parameter m). It follows from the definition of this “nonforking product” that $Y = \{b \in \bar{M} : \mu(\phi(x, b, m)) > 0\}$ is a Borel set over M, m of positive λ_y -measure. By regularity of λ (as a Borel measure on $S_y(M, m)$) there is a closed over M, m set Z say, with $Z \subseteq Y$ and $\lambda(Z) > 0$. By compactness let $b \in Z$ be weakly random for $\lambda|(M, m)$. As $b \in Z \subseteq Y$, $\mu(\phi(x, b, m)) > 0$. As μ is finitely satisfiable in M , there is $a' \in M$ such that $\models \phi(a', b, m)$. By choice of b , $\lambda(\phi(a', y, m)) > 0$, so by finite satisfiability of λ in M there is $b' \in M$ such that $\models \phi(a', b', m)$. This completes the proof. □

In general, the product of measures is not commutative; a measure need not even commute with itself : we can have $\mu_x \otimes \mu_y \neq \mu_y \otimes \mu_x$. The question of commutativity will become central later on. We note at this point that $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$ iff the Borel measure-zero sets of these two measures coincide. This will not be explicitly used in the body of the paper. In the lemma below we take the point of view of a global Keisler measure as a regular probability measure on the relevant Stone space of global types.

Lemma 3.1.7. (NIP) Let μ_x, λ_y be global measures, invariant over some small set.

1. For any definable set $\phi(x, y)$ there is a Borel subset U_ϕ of the space $S_x(\bar{M}) \times S_y(\bar{M})$ (so in the σ -algebra on $S_{xy}(\bar{M})$ generated by rectangles $D_x \times E_y$) such that $\phi(x, y), U_\phi$ are equal up to $\mu_x \otimes \lambda_y$ -measure zero.
2. Commutativity can be checked at the level of the Borel measure-zero ideal : if $\mu_x \otimes \lambda_y(U) = 0$ for any closed U such that $\lambda_y \otimes \mu_x(U) = 0$, then $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$.

Proof. We may write $\mu = \int_{a \in X} p^a, \lambda = \int_{b \in Y} q^b$, where X, Y are the Stone spaces of the Boolean algebra of global definable sets, modulo the measure zero sets of μ, λ respectively; made into measure spaces using the measures induced from μ, λ ; where for $a \in X, b \in Y, p^a, q^b$ are the corresponding invariant types.

1. Given a formula $\phi(x, y)$, let $U_\phi = \{(a, b) \in X \times Y : \phi \in p^a \otimes q^b\}$. We will show below that U_ϕ is Borel up to a measure zero set. Clearly the stated equality holds.
2. The assumption extends from closed to Borel sets : if U' is any Borel set with $\mu_x \otimes \lambda_y(U') = 0$, then $\mu_x \otimes \lambda_y(U) = 0$ for all closed $U \subseteq U'$, so by assumption $\lambda_y \otimes \mu_x(U) = 0$ for all such U , and since this measure is regular, $\lambda_y \otimes \mu_x(U') = 0$. Now let $\phi(x, y)$ be any formula. By (1) there exists a Borel U with $\mu_x \otimes \lambda_y(\phi \triangle U) = 0$, so $\lambda_y \otimes \mu_x(\phi \triangle U) = 0$. But $\mu_x \otimes \lambda_y(U) = \lambda_y \otimes \mu_x(U)$. So $\mu_x \otimes \lambda_y(\phi) = \lambda_y \otimes \mu_x(\phi)$.

To show that U_ϕ is Borel up to measure 0, choose a finite set \mathcal{L}_m of formulas $\phi(x, e)$ such that for any parameter c , there exists a definable set $D \in \mathcal{L}_m$ with $\mu(\phi(x, c) \triangle D) < 2^{-m}$. Let $\mathcal{L} = \cup_m \mathcal{L}_m$, and fix some enumeration of \mathcal{L} (or just of each \mathcal{L}_m). All formulas of \mathcal{L} are defined over some small model M_0 , such that μ, λ are M_0 -invariant.

Any $b \in Y$ determines a weakly random type for λ over M_0 , $q^b|_{M_0}$. Since μ is M_0 -invariant, for $c, c' \models q^b|_{M_0}$ and $D \in \mathcal{L}(M)$ we have $\mu(\phi(x, c) \triangle D) < 2^{-m}$ iff $\mu(\phi(x, c') \triangle D) < 2^{-m}$; so $\mu(\phi(x, c) \triangle \phi(x, c')) = 0$. Thus we will write $\phi(x, b)$ to denote the class of any such $\phi(x, c)$, up to μ measure 0.

Let $D_m(b)$ be the least $D \in \mathcal{L}_m$ such that $\mu(\phi(x, b) \triangle D) < 2^{-m}$.

By the usual proof of completeness of $L^1(\mu)$, $\phi(x, b)$ differs by μ -measure zero from the Borel set

$$D(b) = \{x : (\exists m_0)(\forall m \geq m_0)(x \in D_m(a))\}$$

Since μ is a Borel measure,

$$\{(b, m, D) : D \in \mathcal{L}_m, \mu(\phi(x, b) \triangle D) < 2^{-m}\}$$

is Borel, and so the map $(b, m) \mapsto D_m(b)$ is Borel.

So $E = \{(a, b) : (\exists m_0)(\forall m \geq m_0)(\exists D \in \mathcal{L})(D = D_m(b) \text{ and } a \in D)\}$ is also Borel. For all b , $E(a)$ differs from $U_\phi(b)$ by μ -measure zero. This finishes the proof. \square

Definition 3.1.8. (Assume NIP.) Let μ_x be a global Keisler measure which is invariant (i.e. A -invariant for some small A). Then $\mu_{x_1, \dots, x_n}^{(n)}$ is defined (inductively) by $\mu_{x_1}^{(1)} = \mu_{x_1}$, and $\mu_{x_1, \dots, x_n, x_{n+1}}^{(n+1)} = \mu_{x_{n+1}} \otimes \mu_{x_1, \dots, x_n}^{(n)}$. We put $\mu_{x_1, x_2, \dots}^\omega$ to be the union.

The following notation will be used repeatedly in the paper : If X is a space (usually of types) and $Y \subset X$, then $\text{Fr}(Y; p_1, \dots, p_k)$ denotes $|\{i : p_i \in Y\}|/k$. Similarly the notation $\text{Fr}(\phi(x); a_1, \dots, a_k)$ stands for $|\{i : a_i \models \phi(x)\}|$.

Finally we recall the *weak law of large numbers* in the form we will use it. Any basic text on probability theory is a reference.

Fact 3.1.9. *Let μ be a Borel probability measure on a space X . Let μ^k be the product measure on X^k . Let Y be a measurable subset of X . THEN for any $\epsilon > 0$, $\mu^k(\{(p_1, \dots, p_k) : |\text{Fr}(Y, \bar{p}) - \mu(Y)| < \epsilon\}) \rightarrow 1$ as $k \rightarrow \infty$.*

Thanks to the Wrocław model theory group, in particular H. Petrykowski, for pointing out some errors in an early version of [33], which we deal with in section 5 of the current paper. Some of the results in sections 2 and 3 of the present paper appear in the third authors Master's Thesis [60]. However we do not follow the “formal points” formalism from there.

§3.2 Smooth measures and indiscernibles

Here we discuss smooth measures, using and repeating some material from Keisler's paper [37], but also applying the results in the context of NIP theories to obtain useful results about arbitrary measures as well as “indiscernible measures”.

We will NOT make a blanket assumption that T has NIP.

Definition 3.2.1. A global Keisler measure μ_x is said to be smooth if μ is the unique global extension of $\mu|_M$ for some small model M . We may also call μ smooth over M , and also call $\mu|_M$ smooth.

We should mention that Keisler's notion of a smooth measure was somewhat weaker. He called a Keisler measure over a small set (or even a fragment) if it had a unique global extension modulo the “stable part”. Possibly “minimal” might be a better expression for us, but we stick with our Definition above. A key result of Keisler is Theorem 3.16 from [37] :

Lemma 3.2.2. *(Assume T has NIP.) If μ_x is a Keisler measure over M then it has an extension to a smooth global Keisler measure.*

Note that a complete type (over M , or \bar{M}) is smooth iff it is realized (in M , \bar{M} respectively). A key point of the current paper (also implicit in [33]), is that in an NIP context, one can usefully view a smooth extension of μ as a “realization” of μ , and thus deal effectively with technical issues around measures.

Lemma 3.2.3. *Suppose μ_x is smooth over M . Let $\phi(x, y) \in L$ and $\epsilon > 0$. Then there is some n , formulas $v_i^1(x), v_i^2(x)$ for $i = 1, \dots, n$, and $\psi_i(y)$ for $i = 1, \dots, n$, all over M such that*

- (i) the formulas $\psi_i(\mathbf{y})$ partition \mathbf{y} -space,
- (ii) for all i , if $\models \psi_i(\mathbf{b})$, then $\models \mathbf{v}_i^1(\mathbf{x}) \rightarrow \phi(\mathbf{x}, \mathbf{b}) \rightarrow \mathbf{v}_i^2(\mathbf{x})$, and
- (iii) for each i , $\mu(\mathbf{v}_i^2(\mathbf{x})) - \mu(\mathbf{v}_i^1(\mathbf{x})) < \epsilon$.

Proof. By smoothness of μ and Lemma 1.3 (iv) of [37] for example, for each $\mathbf{b} \in \bar{M}$ there are formulas $\mathbf{v}^1(\mathbf{x})$, $\mathbf{v}^2(\mathbf{x})$ over M , such that

(*) $\models \mathbf{v}^1(\mathbf{x}) \rightarrow \phi(\mathbf{x}, \mathbf{b}) \rightarrow \mathbf{v}^2(\mathbf{x})$, and $\mu(\mathbf{v}^2(\mathbf{x})) - \mu(\mathbf{v}^1(\mathbf{x})) < \epsilon$.

By compactness, there are finitely many such pairs, say, $(\mathbf{v}_i^1(\mathbf{x}), \mathbf{v}_i^2(\mathbf{x}))$ such that for every \mathbf{b} one of these pairs satisfies (*). It is then easy to find the $\psi_i(\mathbf{y})$. \square

Note that it follows from Lemma 2.3 that if μ is a global smooth Keisler measure, then μ is smooth over some model M_0 of cardinality at most $|T|$. Note also that Lemma 2.3 yields a direct way of seeing both the definability over M and finite satisfiability in M of μ .

Definition 3.2.4. If μ_x and λ_y are both Keisler measures over M (with x, y disjoint tuples of variables), then a Keisler measure $\omega_{x,y}$ over M extending both μ_x and λ_y is said to be a *separated amalgam* of μ_x and λ_y , if for any formulas $\phi(x), \psi(y)$ over M , $\omega(\phi(x) \wedge \psi(y)) = \mu(\phi(x)) \cdot \lambda(\psi(y))$.

This is the same thing as saying that $\omega_{x,y}$, as a regular Borel probability measure on $S_{xy}(M)$ extends the product measure $\mu_x \times \lambda_y$ on the space $S_x(M) \times S_y(M)$. Note that as soon as μ_x is not a complete type, there will be at least two extensions of $\mu_x \cup \mu_y$ to Keisler measures over M ; one giving $x = y$ measure 1, which will not be separated, and one extending the product $\mu_x \times \mu_y$ which will be separated. On the other hand if μ_x is a measure over M and $q(y)$ a complete type over M then any amalgam ω_{xy} of μ and q will be separated.

We now give several corollaries of Lemma 2.3.

Corollary 3.2.5. *Suppose μ_x is a smooth global Keisler measure. Then for any global Keisler measure λ_y , there is a unique separated amalgam of μ_x and λ_y .*

Proof. Assume μ is smooth over M . Let $\omega_{x,y}$ be such an amalgam. Let $\phi(x, y) \in L$ and $\epsilon > 0$. Let $\mathbf{v}_i^1(\mathbf{x})$, $\mathbf{v}_i^2(\mathbf{x})$, $\psi_i(\mathbf{y})$, for $i = 1, \dots, n$ be as given by Lemma 2.3. Then for each i , $\models \mathbf{v}_i^1(\mathbf{x}) \wedge \psi_i(\mathbf{y}) \rightarrow \phi(\mathbf{x}, \mathbf{y}) \wedge \psi_i(\mathbf{y}) \rightarrow \mathbf{v}_i^2(\mathbf{x}) \wedge \psi_i(\mathbf{y})$. Let $r_i = \mu(\mathbf{v}_i^1(\mathbf{x}))$ and $t_i = \lambda(\psi_i(\mathbf{y}))$. It follows from the assumptions that $\sum_i r_i t_i \leq \omega_{xy}(\phi(x, y)) \leq \sum_i (r_i + \epsilon) t_i = \sum_i r_i t_i + \epsilon$. Hence $\omega_{xy}(\phi(x, y))$ is determined. \square

Note in particular that a smooth measure μ_x has a unique amalgam with any complete type. Note also that if μ_x is a Borel definable global measure, and λ_y arbitrary then $\mu_x \otimes \lambda_y$ is a separated amalgam. It follows that if μ_x is smooth and λ_y is Borel definable (in the NIP case, invariant) then $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$. Namely a smooth measure “commutes” with any other invariant measure.

Corollary 3.2.6. *Suppose μ_x is smooth over M . Let $\phi(x, y) \in L$ and X_1, \dots, X_k a finite collection of Borel over M sets. Then for any $\epsilon > 0$, for all sufficiently large m , there is a formula $\theta_m(x_1, \dots, x_m)$ such that $\lim_{m \rightarrow +\infty} \mu^{(m)}(\theta_m) = 1$ and for any $(a_1, \dots, a_m) \models \theta_m$,*

(i) for each $b \in M$, $\mu(\phi(x, b))$ is within ϵ of $\text{Fr}(\phi(x, b), a_1, \dots, a_m)$.

Also, we can find such (a_1, \dots, a_m) such that furthermore :

(ii) $\mu(X_i)$ is within ϵ of $\text{Fr}(X_i, a_1, \dots, a_m)$ for each i , and

(iii) for each $b \in \bar{M}$, $\mu(X_i \cap \phi(x, b))$ is within ϵ of $\text{Fr}(X_i \cap \phi(x, b), a_1, \dots, a_m)$.

Proof. Note first that we can assume $x = x$ is an instance of $\phi(x, y)$ and that $x = x$ is included among the X_i . Hence (iii) implies (i) and (ii).

Let $v_i^1(x), v_i^2(x)$ (over M) for $i = 1, \dots, n$ say be given by Lemma 2.3 for $\phi(x, y)$ and $\epsilon/4$. Consider the formula $\theta_m(x_1, \dots, x_m)$ that expresses that $\text{Fr}(v_i^j(x), x_1, \dots, x_m)$ is within $\epsilon/4$ of $\mu(v_i^j(x))$ for each i and j . The weak law of large numbers, Fact 1.8, applied to $X = S_x(M)$, $\mu_x|_M$ (as a probability measure on X), to the Borel sets $v_i^j(x)$ (all i, j), and $\epsilon/4$ implies that $\lim_{m \rightarrow +\infty} \mu^{(m)}(\theta_m(x_1, \dots, x_m)) = 1$. (We here use the fact that $\mu_{(x_1, \dots, x_m)}^{(m)}$ is a separated amalgam of $\mu_{x_1}, \dots, \mu_{x_m}$.)

Next apply the weak law of large numbers, this time to the Borels $v_i^j \cap X_r$ (all i, j, r) and $\epsilon/4$ to obtain suitable types $p_1, \dots, p_m \in S_x(M)$. Let a_1, \dots, a_m be realizations of p_1, \dots, p_m respectively. Let λ_x be the average of the $\text{tp}(a_i/M)$. Let $b \in \bar{M}$, and i be such that $v_i^1(x) \rightarrow \phi(x, b) \rightarrow v_i^2(x)$. Then also for each r , $v_i^1(x) \wedge x \in X_r \rightarrow \phi(x, b) \wedge x \in X_r \rightarrow v_i^2(x) \wedge x \in X_r$. Also clearly $|(\mu(v_i^1(x) \cap X_r) - \mu(v_i^2(x) \cap X_r))| < \epsilon/4$.

Now $\lambda(v_i^j(x) \cap X_r)$ is within $\epsilon/4$ of $\mu(v_i^j(x) \cap X_r)$ for $j = 1, 2$ from which it follows that $\lambda(\phi(x, b) \cap X_r)$ is within ϵ of $\mu(\phi(x, b) \cap X_r)$ giving (iii) (so also (i) and (ii), for (i); note that only the fact that $(a_1, \dots, a_m) \models \theta_m$ is needed). \square

Definition 3.2.7. We will call a global Borel definable Keisler measure μ_x *fim* (a “frequency interpretation measure”) if : for every $\phi(x, y) \in L$, and $\epsilon > 0$, for arbitrary sufficiently large m , there is $\theta_m(x_1, \dots, x_m)$ (with parameters) such that :

(i) $\lim_{m \rightarrow +\infty} \mu^{(m)}(\theta_m) = 1$,

(ii) for all $(a_1, \dots, a_m) \models \theta_m(x_1, \dots, x_m)$, $\mu(\phi(x, b))$ is within ϵ of $\text{Fr}(\phi(x, b), a_1, \dots, a_m)$.

So we have seen that smooth measures are *fim*.

Note that it follows, as in Corollary 2.6, that for any *fim* measure, any $\phi(x, y)$ and any Borel set X , we can find (a_1, \dots, a_m) such that (i) and (ii) of Corollary 2.6 hold. We will see later on that we can also have (iii).

Corollary 2.6 plus Lemma 2.2 enables us to directly prove something about arbitrary measures, for which in [33] we used the Vapnik-Chervonenkis Theorem.

Corollary 3.2.8. *(Assume T has NIP.) Let μ_x be any measure over M . Let $\phi(x, y) \in L$, $\epsilon > 0$, and let X_1, \dots, X_k be Borel sets over M . THEN for all large enough n there are a_1, \dots, a_n such that for all $r = 1, \dots, k$ and all $b \in M$, $\mu(X_r \cap \phi(x, b))$ is within ϵ of $\text{Fr}(X_r \cap \phi(x, b), a_1, \dots, a_n)$.*

Proof. By Lemma 2.2, let μ' be global extension of μ which is smooth over some $M' > M$. Apply Corollary 2.6 to μ' and M' . \square

We will often use the following consequence of this corollary (in a NIP context) : If μ, λ are two global invariant measures, assume that μ commutes with every type p weakly random for λ , then μ and λ commute.

To see this, given a formula $\phi(x, y)$, write $(\mu(x) \otimes \lambda(y))(\phi(x, y)) = \int f(y) d\lambda_y$ as in the paragraph before Lemma 1.5, and approximate that integral by some finite sum $F_N = \sum_{j=1, \dots, N} \lambda(Y_j) c_j$. Use the corollary to find types p_1, \dots, p_n weakly random for λ such that if $a_i \models p_i$ for all i , $\text{Fr}(Y_j \cap \phi(x, b), a_1, \dots, a_n)$ is within ϵ of $\lambda(Y_j \cap \phi(x, b))$ for all j and $b \in \bar{M}$.

If $\tilde{\lambda}$ denotes the average of the types p_1, \dots, p_n then we leave it to the reader to check that $(\mu(x) \otimes \tilde{\lambda}(y))(\phi(x, y))$ is close to $(\mu(x) \otimes \lambda(y))(\phi(x, y))$ and $(\tilde{\lambda}(y) \otimes \mu(x))(\phi(x, y))$ is close to $(\lambda(y) \otimes \mu(x))(\phi(x, y))$. This is enough.

We now begin to discuss “indiscernibles” in the Keisler measure context.

Definition 3.2.9. Let $\mu_{x_1, x_2, \dots}$ be a Keisler measure over M , where x_1, x_2, \dots are distinct variables of the same sort.

- (i) We say that $\mu_{(x_i)_i}$ is indiscernible if for every formula $\phi(x_1, \dots, x_n)$ over M and all $i_1 < i_2 < \dots < i_n$, $\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{i_1}, \dots, x_{i_n}))$.
- (ii) We say that μ is totally indiscernible if $\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{i_1}, \dots, x_{i_n}))$, whenever i_1, \dots, i_n are distinct.

So indiscernibility of a measure μ is with respect to a given sequence $(x_i)_i$ of variables. Likewise we can speak of an indiscernible measure in variables $(x_i : i \in I)$ where I is a totally ordered index set. We can use compactness to “stretch” indiscernible measures in the obvious manner.

We do not know any elementary proof of the following lemma, so we refer the reader to [71]. The lemma is an equivalent formulation of Theorem 5.3 of that paper, the equivalence being a consequence of Lemma 5.4 there.

Lemma 3.2.10. (Assume T has NIP.) If $\mu_{(x_i : i < \omega)}$ over M is indiscernible, $\nu(y)$ is a measure over M , and ω is an amalgam of these over M , and $\phi(x, y) \in L$, then $\lim_{i \rightarrow \infty} \omega(\phi(x_i, y))$ exists. Equivalently, for any such ϕ, μ, ν , and ω , and ϵ , it is not the case that $|\omega(\phi(x_i, y)) - \omega(\phi(x_{i+1}, y))| > \epsilon$ for all i .

Remark 3.2.11. (NIP) In particular, given indiscernible measure $\mu_{(x_i)_i}$ and an extension μ' over M' then we obtain a unique Keisler measure in single variable x over M' , which we call $\text{Av}(\mu', M')$, whose measure of $\phi(x, c)$ (for $c \in M'$) is $\lim_{i \rightarrow \infty} (\mu'(\phi(x_i, c)))$.

Corollary 3.2.12. (NIP) For any formula $\phi(x, y)$ and ϵ there is N such that for any indiscernible Keisler measure $\mu_{(x_i : i < \omega)}$ over a model M (or set A), $b \in \bar{M}$ and extension μ' of μ over (M, b) , there do not exist $i_1 < i_2 < \dots < i_N$ such that $|\mu'(\phi(x_{i_j}, b)) - \mu'(\phi(x_{i_{j+1}}, b))| > \epsilon$ for all $j = 1, \dots, N - 1$.

Proof. By compactness, in the space of measures over M in variables $((x_i)_{i < \omega}, y)$. \square

Corollary 3.2.13. (NIP) Given $\phi(x, y)$ and $\epsilon > 0$ there is N such that for any totally indiscernible Keisler measure $\mu_{(x_i: i < \omega)}$ over a model M , and $b \in \bar{M}$ and extension μ' of μ over (M, b) , for any $r \in [0, 1]$ either $\{i : \mu'(\phi(x_i, b)) \geq r + \epsilon\}$ has cardinality $< N$, or $\{i : \mu'(\phi(x_i, b)) \leq r - \epsilon\}$ has cardinality $< N$.

Proof. Let N be as given by Corollary 2.12 for $\phi(x, y)$ and ϵ . Then clearly it also works for the present result. \square

Lemma 3.2.14. (NIP.) Let μ be a global invariant Keisler measure. Then

(i) $\mu^{(\omega)}$ is indiscernible,

(ii) if both μ and ν are A -invariant and $\mu^{(\omega)}|A = \nu^{(\omega)}|A$ then $\mu = \nu$.

Proof. (i) Obvious and easily proved by induction.

(ii) This is as in the type case : namely suppose for a contradiction that $\mu(\phi(x, b)) = r \neq s = \nu(\phi(x, b))$. Let $\lambda_1(x_1) = \mu(x_1)$, and for even n , $\lambda_n(x_1, \dots, x_n) = \nu(x_n) \otimes \lambda_{n-1}(x_1, \dots, x_{n-1})$ and for odd n , $\lambda_n = \mu(x_n) \otimes \lambda_{n-1}(x_1, \dots, x_{n-1})$. Let $\lambda_{x_1, x_2, \dots}$ be the union. Then one checks that $\lambda|A = \mu^{(\omega)}|A = \nu^{(\omega)}|A$. But $\lambda(\phi(x_i, b)) = r$ for odd i and equals s for even i , contradicting Lemma 2.10. \square

§3.3 Generically stable measures

This section contains our main results. Namely Theorem 3.2 below which gives equivalent conditions for a measure to be generically stable in an NIP theory. We assume NIP throughout, although it would not be uninteresting to develop the theory for arbitrary T . We give two proofs of that theorem, the first one follows very closely the proof of the analogous result for types (Proposition 3.2 of [33]) while the second one is inspired by the proof of the Vapnik-Chervonenkis theorem and avoids the use of Lemma 2.10.

We begin by giving a promised generalization of Lemma 3.4 of [33] to measures.

Lemma 3.3.1. Suppose that μ_x and λ_y are global Keisler measures such that μ is finitely satisfiable (in some small model) and λ is definable. Then $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$.

Proof. We first note that using the remark following 2.8 it suffices to prove the lemma when μ is a type p say. So assume M is a small model over which the TYPE $p(x)$ is finitely satisfiable and the measure λ_y is definable. Let $\phi(x, y) \in L$. Suppose for a contradiction that $\phi(x, y) \in L$ and $(p(x) \otimes \lambda_y)(\phi(x, y)) = r$, $(\lambda_y \otimes p(x))(\phi(x, y)) = s$, and $r \neq s$. Let $\epsilon = |r - s|/4$. Note that s is precisely $\lambda(\phi(a, y))$ where a is some (any) realization of $p|M$. By definability of λ over M , let $\theta(x)$ be a formula over M , which is in $p|M$, and such that

(I) if $\models \theta(a')$ then $\lambda(\phi(a', y))$ is strictly within ϵ of s .

By Borel definability of p over M , $X = \{b : \phi(x, b) \in p\}$ is Borel over M . Hence $r = \lambda(X)$. Apply 2.8 to $\lambda|M$ to find $b_1, \dots, b_n \in M$ such that

(II) for all $a' \in M$, $\lambda(\phi(a', y))$ is within ϵ of $\text{Fr}(\phi(a', y), b_1, \dots, b_n)$ and

(III) the proportion of b_i 's in X is within ϵ of r . Suppose for simplicity that that $b_i \in X$ for precisely $i = 1, \dots, m$.

Let \mathbf{a} realize $\mathbf{p} \mid (M, b_1, \dots, b_n)$. Hence, by the definition of X , $\models \phi(\mathbf{a}, b_i)$ just if $1 \leq i \leq m$. Also of course $\models \theta(\mathbf{a})$. By finite satisfiability of \mathbf{p} in M there is $\mathbf{a}' \in M$ such that $\models \theta(\mathbf{a}')$, $\models \phi(\mathbf{a}', b_i)$ for $i = 1, \dots, m$, and $\models \neg \phi(\mathbf{a}', b_i)$ for $i = m+1, \dots, n$. By (II) $\lambda(\phi(\mathbf{a}', y))$ is within ϵ of m/n . On the other hand by (I) $\lambda(\phi(\mathbf{a}', y))$ is within ϵ of s . As m/n is within ϵ of r , we have a contradiction. \square

Theorem 3.3.2. *Suppose that $\mu(x)$ is a global Keisler measure which is A -invariant. Then the following are equivalent :*

- (i) μ is both definable (necessarily over A) and finitely satisfiable in a small model (necessarily in any model containing A),
- (ii) $\mu_{(x_1, x_2, \dots)}^{(\omega)} \mid A$ is totally indiscernible,
- (iii) μ is fim ,
- (iv) for any global A -invariant Keisler measure λ_y , $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$,
- (v) μ commutes with itself : $\mu_x \otimes \mu_y = \mu_y \otimes \mu_x$.
- (vi) for some small model M_0 containing A , for any Borel over A set X and any formula $\phi(x)$ over \bar{M} , if $\mu(X \cap \phi(x)) > 0$ then there is $\mathbf{a} \in M_0$ such that $\mathbf{a} \in X$ and $\phi(\mathbf{a})$.

Proof. (i) implies (ii) : Note that if μ and λ are both definable measures, then so is $\mu \otimes \lambda$. So we see that each $\mu_{x_1, \dots, x_n}^{(n)}$ is definable. By Lemma 3.1 $\mu_{x_1, \dots, x_{n+1}}^{(n+1)} =_{\text{def}} \mu_{x_{n+1}} \otimes \mu_{x_1, \dots, x_n}^{(n)} = \mu_{x_1, \dots, x_n}^{(n)} \otimes \mu_{x_{n+1}}$. It follows easily (using indiscernibility of each $\mu^{(n)}$) that each $\mu_{x_1, \dots, x_{(n)}}^{(n)}$ is totally indiscernible, hence so is $\mu^{(\omega)}$.

(iii) implies (i) : For all $\phi(x, y) \in L$, $\epsilon = 1/n$ and sufficiently large m , the definition of fim supplies us with a formula $\theta(x_1, \dots, x_m)$. Take a model M containing the parameters of those formulas for all ϕ , n and m . Then μ is definable and finitely satisfiable over M .

(ii) implies (iii) : Without loss $A = M$ is a small model. Assume $\mu^{(\omega)}$ totally indiscernible. *Claim.* Suppose $\lambda_{x_1, x_2, \dots}$ is an extension of $\mu_{x_1, x_2, \dots}^{(\omega)} \mid M$ to a model $M' > M$. Then $\text{Av}(\lambda, M')$ is precisely $\mu \mid M'$.

Proof of Claim. Otherwise, we have some formula $\phi(x, c)$ over M' such that $\mu(\phi(x, c)) = r$ say, and $\text{Av}(\lambda, M')(\phi(x, c)) = s \neq r$. Without loss $s > r$, and let $\epsilon = s - r$. Let $N_{\phi, \epsilon}$ be given by Corollary 2.13. By Lemma 2.10, $\lambda(\phi(x_i, c)) > s - \epsilon$ for eventually all i . However let $\alpha_{(x_i: i < \omega + \omega)}$ be $\mu_{(x_i: \omega \leq i < \omega + \omega)}^{(\omega)} \mid M' \otimes \lambda_{(x_i: i < \omega)}$, a measure over M' . Then α is clearly an extension of $\mu^{(\omega + \omega)}$ to M' . But $\alpha(\phi(x_i, c)) = r$ for all $i \geq \omega$, and we have a contradiction to the existence of $N_{\phi, \epsilon}$. The claim is proved.

Let $\epsilon > 0$ and $\phi(x, y) \in L$. Let N be given by Corollary 2.13 for ϕ and ϵ and let $M = 4N$. Then, by the Claim, for any two measures λ, λ' extending $\mu^{(\omega)}$, for any $b \in \bar{M}$, we have $\{i : |\lambda(\phi(x_i, b)) - \lambda'(\phi(x_i, b))| \geq 2\epsilon\} < 2N$. By compactness, there is a formula $\Phi(x_1, \dots, x_M)$ and a small $r > 0$ such that for any measure $\nu_{(x_1, \dots, x_M)}$ such that $|\nu(\Phi) - \mu^{(M)}(\Phi)| \leq r$, we have $\{i : |\nu(\phi(x_i, b)) - \mu^{(M)}(\phi(x_i, b))| \geq 2\epsilon\} < 2N$.

For a sufficiently large integer k , let $\theta_{kM}(x_1, \dots, x_{kM})$ be the formula which expresses that $\text{Fr}(\Phi(x_1, \dots, x_M); y_0, \dots, y_{k-1})$ is within $r/2$ of $\mu_{(M)}(\Phi)$, where y_i denotes the tuple of variables $(x_{Mi+1}, \dots, x_{M(i+1)})$. If $kM < n < (k+1)M$, define $\theta_n(x_1, \dots, x_n) = \theta_{kM}(x_1, \dots, x_{kM})$. Then by the weak law of large numbers, $\lim_{n \rightarrow +\infty} \mu^{(w)}(\theta_n) = 1$. Furthermore, for sufficiently large n , if $(a_1, \dots, a_n) \models \theta_n$ and ν is the average of $\text{tp}(a_1/\bar{M}), \dots, \text{tp}(a_n/\bar{M})$, then $\nu(\phi(x, b))$ is within 3ϵ of $\mu(\phi(x, b))$ for every $b \in M$. This shows that μ is fim.

(iii) implies (iv). We have to prove that any fim measure commutes with any invariant measure. As above it suffices to prove that an fim measure μ_x commutes with any invariant type.

Let $q(y)$ be such. Assume both μ and q are M -invariant. Let $\phi(x, y) \in L$. Note that $(\mu_x \otimes q(y))(\phi(x, y)) = \mu(\phi(x, b)) = r$ for some (any) b realizing $q|_M$. And also $(q(y) \otimes \mu_x)(\phi(x, y)) = \mu(X) = s$ where $X = \{a : \phi(a, y) \in q(y)\}$ (a Borel set over M). For given ϵ choose a set a_1, \dots, a_k witnessing fim for μ with respect to $\phi(x, y)$ and such that $\text{Fr}(X; a_1, \dots, a_k)$ is within ϵ of $\mu(X)$. Let b realize $q|(M, a_1, \dots, a_k)$. So $\mu(X)$ is within ϵ of $\text{Fr}(\phi(x, b); a_1, \dots, a_k)$. But the latter is within ϵ of $\mu(\phi(x, b))$. So $r = s$ and we are finished.

(iv) implies (v). Obvious.

(v) implies (ii). This follows from associativity of \otimes : for any $k < n$ we have assuming (v), $\mu_{x_1} \otimes \dots \otimes \mu_{x_k} \otimes \mu_{x_{k+1}} \otimes \dots \otimes \mu_{x_n} = \mu_{x_1} \otimes \dots \otimes \mu_{x_{k+1}} \otimes \mu_{x_k} \otimes \dots \otimes \mu_{x_n}$. This is enough.

(vi) is a form of "Borel satisfiability". It is analogous to (i) since Borel definability is automatic. The proof of equivalence with the other conditions is postponed to Lemma 3.6 and Theorem 4.8 below.

□

We call a global Keisler measure μ *generically stable* if it satisfies the equivalent conditions of Theorem 3.2. (Of course assuming T is NIP.)

Proposition 3.3.3. *Suppose that μ is generically stable and A -invariant. Then μ is the unique A -invariant extension of $\mu|_A$.*

Proof. Suppose that ν is A -invariant and $\nu|_A = \mu|_A$. By property (iv) above we check inductively that $\mu^{(n)}|_A = \nu^{(n)}|_A$ for all n . By Lemma 2.14, $\mu = \nu$. □

We give now a proof of Theorem 3.2 which does not use Lemma 2.10. That lemma was used only in the implication (ii) \rightarrow (iii). So we give an alternative proof that if an invariant measure μ is such that $\mu^{(w)}$ is totally indiscernible, then μ is fim.

By a symmetric measure on some X^n , X a definable set, we will now mean a measure $\mu_{(x_1, \dots, x_n)}$ such that $\mu(x_i \in X) = 1$ for all i and for any $\sigma \in S_n$ and formula $\phi(x_1, \dots, x_n)$, $\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{\sigma,1}, \dots, x_{\sigma,n}))$.

The following crucial lemma is related to the classical Vapnik-Chervonenkis theorem (see [68]) and could be proved by similar methods. But it does not seem to be a direct consequence of it.

Lemma 3.3.4. *Let $\phi(x, y) \subseteq X \times Y$ be a formula over a model M . For $n > 0$, let μ_n be any symmetric, M -invariant global measure on X^{2n} . Given $b \in Y$ and $a = (a_1, \dots, a_n) \in X^n$, let $f(a; b) = \text{Fr}(\phi(x, b); a_1, \dots, a_n)$. Let*

$$\delta_0(a, a'; b) = |f(a; b) - f(a'; b)|,$$

$$\delta(a, a') = \sup_{b \in M} \delta_0(a, a'; b).$$

Finally, let $E(n)$ be the μ_n -expectation of δ . Then $\lim_{n \rightarrow \infty} E(n) = 0$.

Proof. Note first that δ is measurable for the boolean algebra generated by the definable subsets of X^{2n} of the form $(\exists y)(\bigwedge_{i=1}^{2n} \phi(x_i, y)^{v(i)})$ (where $\phi(x_i, y)^{v(i)}$ is either $\phi(x_i, y)$ or $\neg\phi(x_i, y)$). In particular it makes sense to ask for the μ_n -expectation of δ .

Fix $\epsilon > 0$, and let n be large compared to ϵ .

Let $\mathbb{Z}/2$ act on the variables $\{x_i, x'_i\}$ by flipping them, and let $(\mathbb{Z}/2)^n$ act on the set $\{x_1, \dots, x_n, x'_1, \dots, x'_n\}$ by the product action.

Given (a, a') and b , we have $|\delta_0(a, a'; b)| \leq 1$. Let

$$X(b) = \{s \in (\mathbb{Z}/2)^n : \delta_0(s(a, a'); b) > \epsilon\}$$

and let $c(b) = \{i : \phi(a_i, b) \neq \phi(a'_i, b)\}$. If $c(b) = \emptyset$ then $\delta_0(s(a, a'); b) = 0$ for all s . Otherwise, we view $\delta(s(a, a'); b)$ as a random variable of s (on a finite probability space). More precisely, it is the absolute value of a sum of $|c(b)|$ independent random variables each of expectation 0 and variance $1/n^2$. Therefore $\delta_0(s(a, a'); b)$ has expectation 0, and variance $|c(b)|/n^2$. By Tchebychev's inequality we have $|X(b)| \leq |c(b)|/(n\epsilon)^2 \leq \epsilon^{-2}$. So $|X(b)|/2^n \leq \epsilon^{-2}2^{-n}$.

Now $X(b)$ depends only on $\{i : \phi(a_i, b)\}$ and $\{i : \phi(a'_i, b)\}$; there are polynomially many possibilities for these sets, by NIP. Hence, $\cup_b X(b)$ is an exponentially small subset of $(\mathbb{Z}/2)^n$. If n is large enough, it has proportion $< \epsilon$. Let $s \notin \cup_b X(b)$. Then $|\delta_0(s(a, a'), b)| \leq \epsilon$ for all b . If $s \in \cup_b X(b)$ we have at any rate $|\delta_0(s(a, a'), b)| \leq 1$. Thus $2^{-n} \sum_s \sup_b \delta(s(a, a'); b) < 2\epsilon$.

By the symmetry of μ_n , $E(n)$ equals the μ_n -expectation of $\sup_a \delta(s(a, a'))$ for any $s \in (\mathbb{Z}/2)^n$, hence it is also equal to the average $2^{-n} \sum_s \sup_a \delta(s(a, a'))$. So $E(n) < 2\epsilon$. \square

Corollary 3.3.5. *Let μ be an M -invariant global measure, such that $\mu^{(\omega)}$ is totally indiscernible, then μ is fin.*

Proof. Let $\phi(x, y)$ be a formula, and take $\epsilon > 0$. By the previous lemma, for large enough n , the set $W = \{(a, a') : \delta(a, a') < \epsilon/4\}$ satisfies $\mu^{(2n)}(W) \geq 1 - \epsilon$. Note that this is a definable set. Therefore there exists a such that $\mu^{(n)}(W(a)) \geq 1 - 2\epsilon > 1/2$. (Where $W(a) = \{a' : (a, a') \in W\}$.) Now fix $a \in \bar{M}$. Then for all b' , $\delta_0(a, a'; b) \leq \delta(a, a')$.

On the other hand let $Q'_n(b)$ be the set of a' such that $|f(a'; b) - \mu(\phi(x, b))| \geq \epsilon/2$. By Tchebychev's inequality, and since the variance of the truth value of $\phi(x, b)$ is at most 1, we have $\mu^{(n)}(Q'_n(b)) \leq 1/(n(\epsilon/2)^2)$. Let $Q_n(b)$ be the complement of $Q'_n(b)$, and assume

$n > 2(\epsilon/2)^{-2}$ (note that this does not depend on b). Then $\mu^{(n)}(Q_n(b)) > 1/2$. Hence there exists $a' \in W(a) \cap Q_n(b)$. So $\delta_0(a, a'; b) < \epsilon/4$ and $|f(a'; b) - \mu(\phi(x, b))| < \epsilon/2$. Now for any $a'' \in W(a)$, we have $\delta_0(a', a''; b) < \epsilon/2$, therefore $|f(a'', b) - \mu(\phi(x, b))| < \epsilon$.

So we have found, for large enough n a formula $\theta'_n = W(a)$ satisfying condition (ii) in the definition of fim with $\mu^{(n)}(\theta'_n) \geq 1 - 2\epsilon$. As this is true for all ϵ , we can construct a sequence of formulas $\theta_n(x_1, \dots, x_n)$ satisfying the same condition, but with $\mu^{(n)}(\theta_n) \rightarrow 1$. This proves that μ is fim . \square

Lemma 3.3.6. *Let μ be an M -invariant global fim measure. Let $\phi(x, y)$ be a formula over M and let X be a Borel over M set. Then for any $\epsilon > 0$, for some m , we can find (a_1, \dots, a_m) such that for each $b \in \bar{M}$, $\mu(X \cap \phi(x, b))$ is within ϵ of $\text{Fr}(X \cap \phi(x, b), a_1, \dots, a_m)$.*

Proof. We know that $\mu^{(\omega)}$ is totally indiscernible. The proof is then a slight modification of the lemma above and its corollary. First, in the lemma, change the definition of f to $f(a; b) = \text{Fr}(\phi(x, b) \wedge x \in X; a_1, \dots, a_n)$. Define δ_0 and δ accordingly. Then the proof goes through without any difficulties. The corollary also goes through with the new definitions of f , δ_0 and δ , only W and $W(a)$ are no longer definable. Still, $W(a)$ is a Borel set of measure greater than $1/2$ and for any $a' \in W(a)$, and any b , we have $|f(a', b) - (\mu(\phi(x, b) \cap X))| < \epsilon$. \square

Finally we will point out how generically stable measures are very widespread in NIP theories, in fact can be constructed from any indiscernible sequence. By an indiscernible *segment* we mean $(a_i : i \in [0, 1])$ which is indiscernible with respect to the ordering on the real unit interval $[0, 1]$. For any formula $\phi(x, b)$, $\{i \in [0, 1] : \models \phi(a_i, b)\}$ is a finite union of convex sets, and hence intervals. (See [2] for example.) We define a measure μ_x as follows : $\mu(\phi(x, b))$ is the Lebesgue measure of $\{i \in [0, 1] : \models \phi(a_i, b)\}$ (i.e. just the sum of the lengths of the relevant disjoint intervals). Clearly μ_x is a global Keisler measure on the sort of x .

Proposition 3.3.7. *The global Keisler measure μ_x constructed above is generically stable.*

Proof. Let $A = \{a_i : i \in [0, 1]\}$. We show that μ is both finitely satisfiable in A and definable over A . Finite satisfiability is clear from the definition of μ . (If $\mu(\phi(x, b)) > 0$ then the Lebesgue measure of $C = \{i : \models \phi(a_i, b)\}$ is > 0 hence $C \neq \emptyset$.)

To show definability, we in fact note that μ is fim . First note that for any formula $\phi(x, y)$, there is N_ϕ such that for all b $\{i \in [0, 1] : \models \phi(x, b)\}$ is a union of at most N disjoint intervals. Hence, given $\epsilon > 0$, if we choose $0 < \delta < \epsilon/N$, and let $i_k = 0 + k\delta$ for k such that $k\delta \leq 1$, then for any b , $\mu(\phi(x, b))$ is within ϵ of the proportion of i_k such that $\models \phi(a_{i_k}, b)$. \square

§3.4 The fsg property for groups, types and measures

We will again make a blanket assumption that T has NIP (but it is not always needed). In [32] and [33] definable groups G with finitely satisfiable generics (fsg) played an important role. This fsg property asserted the existence of a global type of G every left translate of which was finitely satisfiable in some given small model. By definition this is a property of a definable group, rather than of some global type or measure. We wanted to find adequate generalizations of the fsg notion to arbitrary complete types $p(x)$ over small sets, and even arbitrary Keisler measures over small sets. A tentative definition of a complete type $p(x) \in S(A)$ having fsg was given in [33]. We try to complete the picture here, making the connection with generically stable global measures. We should say that the subtlety of the fsg notion is really present in the case where the set A is NOT bounded closed. In the group case this corresponds to the case where $G \neq G^{00}$.

We first return to the group case, adding to results from [33].

Recall the original definition :

Definition 3.4.1. The definable group G has fsg if there is a global complete type $p(x)$ of G and a small model M_0 such that every left translate of p is finitely satisfiable in M_0 .

As pointed out in [32] M_0 can be chosen as any model over which G is defined. In [33] the notion was generalized to type-definable groups G .

Remark 3.4.2. The definable group G has fsg if and only if there is a global left invariant Keisler measure μ on G which is finitely satisfiable in some (any) small model M_0 over which G is defined.

Proof. Assuming that G has fsg, Proposition 6.2 of [32] produces the required measure μ . Conversely, supposing that μ is a global left invariant Keisler measure on G , finitely satisfiable in M_0 , let p be some global type of G such that $\mu(\phi) > 0$ for all $\phi \in p$. Namely p comes from an ultrafilter on the Boolean algebra of definable subsets of G modulo the equivalence relation $X \sim Y$ if $\mu(X \Delta Y) = 0$. Then every left translate of every formula in p has μ measure > 0 so is realized in M_0 . \square

The following strengthens the “existence and uniqueness” of (left/right) G -invariant global Keisler measures for fsg groups G , from [33]. But the proof is somewhat simpler, given the results established in earlier sections.

Theorem 3.4.3. *Suppose G has fsg, witnessed by a global left invariant Keisler measure μ finitely satisfiable in small M_0 . Then*

- (i) μ is the unique left invariant global Keisler measure on G , as well as the unique right invariant Keisler measure on G .
- (ii) μ is both the unique left G^{00} -invariant global Keisler measure, as well as the unique right G^{00} -invariant global Keisler measure on G , which extends Haar measure h on G/G^{00} .

Proof. (i) This is precisely 7.7 of [33]. But note that we can use Lemma 3.1 of the present paper and the relatively soft 5.8 of [33] in place of 7.3 and 7.6 of [33]. (Details : Suppose λ is also global left invariant. By the Lemma 5.8 of [33], we may assume λ to be definable. Given definable subset D of G , let $Z = \{(g, h) : g \in hZ\}$. By 3.1, $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$, whence $(\mu_x \otimes \lambda_y)(Z) = \mu(D)$ and $(\lambda_y \otimes \mu_x)(X) = \lambda(D^{-1})$. So $\lambda = \mu^{-1}$. This in particular yields that $\mu = \mu^{-1}$. So $\lambda = \mu$ and μ is also the unique right invariant Keisler measure.) (ii) Note first that μ_x induces a left invariant measure on G/G^{00} which has to be (normalized) Haar measure, and of course μ is (left/right) G^{00} -invariant. Let λ_y be another global left G^{00} -invariant Keisler measure extending (or inducing) Haar measure on G/G^{00} . As in 5.8 of [33] we may assume that λ is definable.

Let X be a definable subset of G . Let $r = \lambda(X)$. Choose $\epsilon > 0$. As G^{00} stabilizes λ , and λ is definable, there is a definable subset Y of G which contains G^{00} and such that for all $g \in Y$, $\lambda(gX) \in (r - \epsilon, r + \epsilon)$. Let $\pi : G \rightarrow G/G^{00}$ be the canonical surjective homomorphism. Then $\{c \in G/G^{00} : \pi^{-1}(c) \subseteq Y\}$ is an open neighbourhood W of the identity in G/G^{00} . Let $U = \pi^{-1}(W) \subseteq Y \subseteq G$. Note that $\mu(U) > 0$ (as it equals the Haar measure of the open subset W of G/G^{00}). We have

(1) for all $g \in U$, $\lambda(gX) \in (r - \epsilon, r + \epsilon)$.

Whence

(2) $(r - \epsilon)\mu(U) \leq \int_{g \in U} \lambda(gX) d\mu \leq (r + \epsilon)\mu(U)$.

Now let $Z = \{(h, g) : h \in gX \text{ and } g \in G\}$. Then

(3) $(\lambda_y \otimes \mu_x)(Z) = \int_{g \in U} \lambda_y(gX) d\mu_x$.

On the other hand clearly

(4) $(\mu_x \otimes \lambda_y)(Z) = \int_{h \in G} \mu_x(hX^{-1} \cap U) d\lambda_y$.

Note that the value of $\mu_x(hX^{-1} \cap U)$ depends only on h/G^{00} , hence as λ and μ agree “on G/G^{00} ”, we see that

(5) $(\mu_x \otimes \lambda_y)(Z) = (\mu_x \otimes \mu_y)(Z)$.

By 3.1 applied twice (to $(\mu_x \otimes \lambda_y)$ AND to $(\mu_x \otimes \mu_y)$) together with (5) we see that

(6) $(\lambda_y \otimes \mu_x)(Z) = (\mu_y \otimes \mu_x)(Z)$.

But $(\mu_y \otimes \mu_x)(Z) = \int_{g \in U} \mu_y(gX) d\mu_x = \mu(X)\mu(U)$.

So using (2) and (3) we see that

(7) $(r - \epsilon)\mu(U) \leq \mu(X)\mu(U) \leq (r + \epsilon)\mu(U)$.

So $r - \epsilon \leq \mu(X) \leq r + \epsilon$. As this is true for all ϵ we conclude that $\mu(X) = r = \lambda(X)$. \square

Remark 3.4.4. The definable group G has fsg if and only if G has a global generically stable left invariant measure.

We now consider the general situation. We first recall the definition from [33] :

Definition 3.4.5. $p(x) \in S(A)$ has fsg if p has a global extension p' such that for any formula $\phi(x) \in p'$ and $|A|^+$ -saturated model M_0 containing A , there is $a \in M_0$ realizing p such that $\models \phi(a)$.

Lemma 3.4.6. $p(x) \in S(A)$ has fsg iff there is a global A -invariant measure μ_x extending p such that whenever $\phi(x)$ is a formula over \bar{M} with μ -measure > 0 , then for any $|A|^+$ -saturated model M_0 containing A , there is $a \in M_0$ realizing p such that $\models \phi(a)$.

Proof. RHS implies LHS. This is trivial because any weakly random type for the measure μ will satisfy Definition 4.5.

LHS implies RHS : Let p' be as given by Definition 4.5. Then 7.12 (i) of [33] says that p' is a nonforking extension of p . Moreover it is clear that any $\text{Aut}(\bar{M}/A)$ -conjugate of p' also satisfies Definition 4.5. Let μ be the global A -invariant measure extending p , constructed from p' in Proposition 4.7 of [33]. Then any formula with positive μ -measure must be in some $\text{Aut}(\bar{M}/A)$ -conjugate of p' so is satisfied in any saturated model M_0 containing A by a realization of p . \square

The last lemma motivates a definition of fsg for arbitrary measures over A .

Definition 3.4.7. Let μ_x be a Keisler measure over A . We say that μ has fsg, if μ has a global A -invariant extension μ' such that for any Borel over A set X , formula $\phi(x)$ over \bar{M} , and $|A|^+$ -saturated model M_0 containing A , if $\mu'(X \cap \phi(x)) > 0$ then there is $a \in M_0$ such that $a \in X$ and $\phi(a)$.

Theorem 3.4.8. Let μ_x be a measure over a set A , then the following are equivalent :

- (i) μ has a unique A -invariant global extension μ' that is moreover generically stable.
- (ii) μ has fsg.
- (iii) μ has a global A -invariant extension μ' such that for any Borel over A set X , formula $\phi(x)$ over \bar{M} , and $|A|^+$ -saturated model M_0 containing A , if $\mu(X) + \mu(\phi(x)) > 1$ then there is $a \in M_0$ such that $a \in X$ and $\phi(a)$.

Proof. (i) implies (ii) : Follows from Lemma 3.3.6.

(ii) implies (iii) is clear.

(iii) implies (i) : We fix global A -invariant measure μ' extending μ and witnessing the assumption. We will prove that μ' commutes with every A -invariant measure. It will follow that $\mu^{(\omega)}|A$ is totally indiscernible, so μ' is generically stable. Uniqueness follows from 3.3. The proof will be a bit like that of 3.1. In fact it is easy to see that μ commutes with any A -invariant type. If $A = \text{bdd}(A)$ this would suffice (as every measure non-forking over A is “approximated” by types non-forking over A using Corollary 3.2.8, which are $\text{bdd}(A)$ -invariant). But for arbitrary A it does not seem to suffice.

So let us fix an A -invariant (thus Borel-definable over A) global measure λ_y .

Let $P = (\mu_x \otimes \lambda_y)(\phi(x, y)) = \int \mu(\phi(x, b)) d\lambda$ and $R = (\lambda_y \otimes \mu_x)(\phi(x, y)) = \int \lambda(\phi(a, y)) d\mu$. We want to show that $P = R$.

For any $t \in [0, 1]$, let $C_t = \{q \in S(A) : \mu(\phi(x, b)) \geq t \text{ for any } b \models q\}$ and $B_t = \{p \in S(A) : \lambda(\phi(a, y)) \geq t \text{ for any } a \models p\}$. These sets are Borel over A .

Let $\epsilon > 0$ and take $N \geq 1/\epsilon$ such that

$$\left| P - \frac{1}{N} \sum_{k=0}^{N-1} \lambda(C_{k/N}) \right| \leq \epsilon.$$

Take a model M containing A and $|A|^+$ -saturated. By Corollary 2.8, there exist n and $p_1, \dots, p_n \in S(M)$ such that $\text{Fr}(\phi(a; y); p_1, \dots, p_n)$ is within ϵ of $\lambda(\phi(a; y))$ for every $a \in M$ and $\text{Fr}(C_{k/N}; p_1, \dots, p_n)$ is within ϵ of $\lambda(C_{k/N})$ for every $k < N$. Realize p_1, \dots, p_n

in \bar{M} by b_1, \dots, b_n respectively. Call $\tilde{\lambda}$ the average measure of b_1, \dots, b_n (seen as global measures).

By construction, we have

$$\left| \frac{1}{N} \sum_{k=0}^{N-1} \lambda(C_{k/N}) - \frac{1}{N} \sum_{k=0}^{N-1} \tilde{\lambda}(C_{k/N}) \right| \leq \epsilon.$$

On the other hand, for all $k < N$:

$$\tilde{\lambda}(C_{k/N}) = \frac{1}{n} |\{i : p_i \in C_{k/N}\}| = \frac{1}{n} |\{i : \mu(\phi(x, b_i)) \geq k/N\}|.$$

It follows that :

$$\left| \frac{1}{N} \sum_{k=0}^{N-1} \tilde{\lambda}(C_{k/N}) - \frac{1}{n} \sum_{i=1}^n \mu(\phi(x, b_i)) \right| \leq \frac{1}{N} \leq \epsilon.$$

Now, for $k \leq N$ let $\Theta_k(x)$ be the formula that says “There are at least k values of i for which $\models \phi(x, b_i)$ holds”. Then (looking at the Venn diagram generated by the sets $\phi(x, b_i)$ and counting each time each region appears in both sums) we see that

$$\frac{1}{n} \sum_{k=1}^n \mu(\Theta_k(x)) = \frac{1}{n} \sum_{i=1}^n \mu(\phi(x, b_i)).$$

Call P' the value of those two sums.

By the construction of $\tilde{\lambda}$, we have the inclusions $B_{k/n+\epsilon}(M) \subseteq \Theta_k(M) \subseteq B_{k/n-\epsilon}(M)$, and fsg for μ implies that $\mu(B_{k/n+\epsilon}) \leq \mu(\Theta_k(x)) \leq \mu(B_{k/n-\epsilon})$.

Then, choosing l such that $l/n \leq 2\epsilon$,

$$\frac{1}{n} \sum_{k=1}^n \mu(B_{k/n+l/n}) \leq P' \leq \frac{1}{n} \sum_{k=1}^n \mu(B_{k/n-l/n}).$$

The difference between the two sums to the right and to the left of P' is at most 8ϵ . We may assume that n was chosen large enough so that $|R - \frac{1}{n} \sum_{k=1}^n \mu(B_{k/n})| \leq \epsilon$. This latter sum satisfies the same double inequality as P' . Therefore $|P' - R| \leq 8\epsilon$. Putting everything together, we see that $|P - R| \leq 11\epsilon$.

As ϵ was arbitrary, we are done. \square

§3.5 Generic compact domination

In [32] the authors introduced the notion of “domination” or control of a type-definable set X by a compact space C equipped with a measure (or ideal) μ : namely there is a “definable” surjective function $\pi : X \rightarrow C$ such that for every (relatively)

definable subset Y of X , for almost all $c \in C$ in the sense of μ , either $\pi^{-1}(c) \subseteq Y$ or $\pi^{-1}(c) \cap Y = \emptyset$. Here of course X, π are defined over a fixed set A of parameters, and Y is definable with arbitrary parameters. There was also a “group version” where $X = G$ is a (type)-definable group, C is a compact group, π a homomorphism, and μ is Haar measure on C .

In this section we will consider a weaker version of compact domination where X is replaced by a suitable space of “generic” types, and we expand on and correct some results which had appeared in a first version of [33]. We view this weak domination as a kind of measure-theoretic weakening of the finite equivalence relation theorem (or even the Open Mapping Theorem). Let us begin by explaining this interpretation. If T is a stable theory, $\text{acl}(A)$ is a small algebraically closed (in \bar{M}^{eq}) set, and $p(x) \in S(\text{acl}(A))$ a type, then the finite equivalence relation theorem says that p has a unique global nonforking extension. Let \mathcal{P} denote the set of global complete types in x which do not fork over A , then \mathcal{P} is a closed subset of $S_x(\bar{M})$ and the restriction map π taking $p(x) \in \mathcal{P}$ to $p|_{\text{acl}(A)}$ is a continuous bijection between the compact spaces \mathcal{P} and $S_x(\text{acl}(A))$, and therefore a homeomorphism. We deduce (still in this stable context) that any Keisler measure μ_x over $\text{bdd}(A)$ has a unique global nonforking ($\text{bdd}(A)$ -invariant) extension : For any formula $\phi(x)$ over \bar{M} there is a formula $\psi(x)$ over $\text{acl}(A)$ such that the image under π of the clopen subset of \mathcal{P} determined by $\phi(x)$ is the clopen subset of $S_x(\text{acl}(A))$ given by $\psi(x)$. Defining $\mu'(\phi(x))$ to be $\mu(\psi(x))$ gives the required unique global nonforking extension of μ .

The question we deal with is the following. Work in a NIP environment. Let $\text{bdd}(A)$ be now the “bounded closure” of a small set A , let \mathcal{P} be the set of global complete types $p(x)$ which do not fork over A (equivalently are $\text{bdd}(A)$ -invariant), C the set of restrictions of types in \mathcal{P} to $\text{bdd}(A)$, and π the restriction map. Both \mathcal{P} and C are naturally, compact spaces and π is continuous, but not in general a homeomorphism. Now suppose that μ_x is a Keisler measure over $\text{bdd}(A)$ with a unique global nonforking extension. What does it imply about $\pi : \mathcal{P} \rightarrow C$? We will prove (5.4 below) that \mathcal{P} is *dominated* by (C, π, μ) , in the sense that for any formula $\phi(x)$ over \bar{M} , the collection of $p \in C$ such that $\pi^{-1}(p)$ intersects both (the clopen set determined by) ϕ and its complement, has μ -measure 0. In fact we prove that domination of \mathcal{P} by (C, π, μ) is *equivalent* to μ having a unique global nonforking extension. We will then give the analogous definable group versions. A global measure on G which is invariant under (left) translation by G^{00} will be the analogue of a global measure which is $\text{bdd}(\emptyset)$ invariant. Haar measure on G/G^{00} will be the analogue of a Keisler measure over $\text{bdd}(\emptyset)$ and the analogue (5.7) of 5.4 will be the statement that Haar measure on G/G^{00} has a unique lifting to a left G^{00} -invariant global Keisler measure on G if and only if the collection of left G^{00} -invariant global types of G is dominated by $(G/G^{00}, \pi, h)$ where π is the natural map and h is Haar measure.

We will again be assuming that T has NIP but it is not always needed.

The first two lemmas are not required for Theorem 5.4, but are worth noting and may be relevant for versions of 5.4 over A rather than $\text{bdd}(A)$.

Lemma 3.5.1. *Let μ_x a Keisler measure over A . Then there is a unique Keisler measure μ'_x over $\text{bdd}(A)$ which extends μ_x and is $\text{Aut}(\text{bdd}(A)/A)$ -invariant.*

Proof. We identify a Keisler measure on $\text{bdd}(A)$ with a regular probability measure on $S_x(\text{bdd}(A))$. Likewise identify μ with a measure on $S(A)$. Now for each $p(x) \in S(A)$, there is a unique A -invariant Keisler measure on $S(\text{bdd}(A))$ extending p , which we call μ_p : the space of extensions of p over $\text{bdd}(A)$ is, as mentioned above, a homogeneous space for the compact Lascar group $\text{Aut}(\text{bdd}(A)/A)$ so has a unique A -invariant measure, which is precisely μ_p . Now define μ' as follows : for a Borel set B over $\text{bdd}(A)$, put $\mu'(B) = \int_{p \in S(A)} \mu_p(B) d\mu$. μ' is clearly A -invariant and we leave to the reader to check uniqueness. \square

Lemma 3.5.2. *Suppose μ_x is a Keisler measure over A which does not fork over A . Then μ has a global A -invariant extension.*

Proof. Let λ_x be some global nonforking extension of μ . By [33], λ is $\text{bdd}(A)$ -invariant and moreover Borel definable over $\text{bdd}(A)$. Fix a formula $\phi(x, b)$. Let Q be the set of complete extensions of $\text{tp}(b/A)$ over $\text{bdd}(A)$. As above Q is a homogeneous space for the compact Lascar group over A and inherits a corresponding measure h say. Define $\mu''(\phi(x, b)) = \int_{b' \in Q} \lambda(\phi(x, b')) dh$. \square

Lemma 3.5.3. *Suppose μ_x is a Keisler measure over A which has a unique global A -invariant extension. Then for any closed subset B of $S_x(A)$ of positive μ -measure, the localization μ_B of μ at B also has a unique global A -invariant extension.*

Proof. Let μ' be the unique global A -invariant extension of μ . Then μ'_B is clearly an A -invariant extension of μ_B . If it is not the unique one, let λ be another A -invariant extension of μ_B . Define a global measure μ'' , by $\mu''(X) = \lambda(X) \cdot \mu(B) + \mu'(X \cap B^c)$ (where B^c is the complement of B in $S(A)$ and X a definable set). Note that μ'' also extends μ and is A -invariant (as the two terms in the sum are A -invariant). However by choosing X such that $\lambda(X) \neq \mu'_B(X)$ we see that $\mu''(X) \neq \mu(X)$, contradicting our assumption. \square

Notation for Theorem 5.4 is as in the third paragraph of this section. A is a small set of parameters, and μ_x a Keisler measure over $\text{bdd}(A)$ (in fact the results will be valid for any boundedly closed set of hyperimaginaries in place of $\text{bdd}(A)$). \mathcal{P} is the set of global complete types $p(x)$ which are $\text{bdd}(A)$ -invariant, equivalently do not fork over A , $C = \{p|_{\text{bdd}(A)} : p \in \mathcal{P}\}$, and $\pi : \mathcal{P} \rightarrow C$ the restriction map. \mathcal{P} and C are closed subspaces of the type spaces $S_x(M)$ and $S_x(\text{bdd}(A))$ respectively, and π is a continuous surjection. μ induces a regular probability measure on C which we also call μ . The only assumption we will make is that \mathcal{P} is nonempty. When appropriate we identify a definable set (of sort x) with a clopen subset of \mathcal{P} .

Theorem 3.5.4. *μ has a unique global nonforking ($\text{bdd}(A)$ -invariant) extension if and only if \mathcal{P} is dominated by (C, π, μ) .*

Proof. Assume the right hand side. We will define a global Keisler measure λ_x extending μ as follows : Let X be a definable set (of sort x). Let $D_0 = \{c \in C : \pi^{-1}(c) \cap X \neq \emptyset \text{ and } \pi^{-1}(c) \cap X^c \neq \emptyset\}$. Then $D_0 \subseteq C$ is closed and $\mu(D_0) = 0$ by the domination assumption. $C \setminus D_0$ is the disjoint union of Borel sets D_1 and D_2 , where $\pi^{-1}(c) \subseteq X$ for $c \in D_1$ and $\pi^{-1}(c) \cap X = \emptyset$ for $c \in D_2$. We define $\lambda(X)$ to be $\mu(D_1)$. It is routine to check both that λ extends μ and that λ is the unique $\text{bdd}(A)$ -invariant extension of μ .

For the converse, assume that μ has a unique global $\text{bdd}(A)$ -invariant extension, say λ . If the domination statement fails, there is a closed subset D of C of positive μ -measure and formula $\phi(x)$ over \bar{M} , such that $\pi^{-1}(c)$ intersects both $\phi(x)$ and $\neg\phi(x)$ for all $c \in D$. Hence for every $p(x) \in D$, both $p(x) \cup \{\phi(x)\}$ and $p(x) \cup \{\neg\phi(x)\}$ do not fork over A . Let μ_D be the localization of μ at D .

Claim. There are ν_1, ν_2 , global nonforking extensions of μ_D such that $\nu_1(\phi(x)) = 1$ and $\nu_2(\phi(x)) = 0$.

Proof of claim. Consider the fragment G generated by the partial types over $\text{bdd}(A)$, $\phi(x)$ and the set Ψ of formulas $\psi(x)$ (over M) which fork over $\text{bdd}(A)$. For each $p \in D$, let $r_p = p(x) \cup \{\phi(x)\} \cup \{\neg\psi(x) : \psi \in \Psi\}$. Then $D' = \{r_p : p \in D\}$ is a closed subset of $S(G)$ and $f : D \rightarrow D'$ defined by $f(p) = r_p$ is a homeomorphism. Using f to define a measure ν'_1 supported on D' gives an extension of μ which assigns 1 to $\phi(x)$. Any extension of ν'_1 to a global Keisler measure ν_1 is a nonforking extension of μ assigning 1 to $\phi(x)$.

Likewise we find μ_2 .

So the claim is proved and gives a contradiction to Lemma 5.3. This completes the proof of the Theorem. □

We leave it as an open problem to find an equivalence to “ μ has a unique global A -invariant extension” where μ is a Keisler measure over an arbitrary (not necessarily boundedly closed) set A . Note that in the stable case any Keisler measure μ over any set A has a unique global A -invariant extension.

Finally in this section we return to definable groups. The relevant uniqueness statement will be something like 4.3(ii). The domination statement will be roughly the domination of a suitable family of “generic” types by G/G^{00} with its Haar measure. We start by tying up a few loose ends.

If p is a global type, the (left) stabilizer of p denoted by $\text{Stab}_l(p)$ is the the set of $g \in G$ such that $g.p = p$ (where of course $g.p$ is the type of ga where $a \models p$).

Lemma 3.5.5. *Let G be a definable group. Suppose there is a global type p of G with $\text{Stab}_l(p) = G^{00}$. Then there is a global left G^{00} -invariant measure μ on G which lifts (extends) Haar measure on G/G^{00}*

Proof. Let h be Haar measure on G/G^{00} which we can think of as a Keisler measure on a suitable fragment F (in fact the fragment consisting of the preimages of closed

sets under $\pi : G \rightarrow G/G^{00}$. Let $p(x)$ be as given by the assumptions. We may assume that p concentrates on G^{00} . Note that $\text{Stab}_1(ap) = G^{00}$ for every translate ap of p . In particular for $c \in G/G^{00}$ there is a unique translate of p by some a in the coset c , so we just write it cp . Note that for each definable subset X of G , and $g \in G^{00}$ we have that $X \triangle gX \notin cp$ for all c . It follows as in the proof of 5.4 that h extends to a Keisler measure μ over the fragment generated by F and all definable sets $X \triangle gX$ ($g \in G^{00}$), such that $\mu(X \triangle gX) = 0$ for all such X, g , and so to a global Keisler measure which is G^{00} -invariant. \square

Lemma 3.5.6. *Suppose μ is a global Keisler measure on G which is (left) G^{00} -invariant. Then $\mu(X \triangle gX) = 0$ for all definable X and $g \in G^{00}$. In particular for all μ -weakly random global p , $\text{Stab}_1(p) = G^{00}$.*

Proof. This is a kind of group version of the fact that if a global Keisler measure is $\text{bdd}(A)$ -invariant (does not fork over $\text{bdd}(A)$) then $\mu(X \triangle X') = 0$ for any $\text{bdd}(A)$ -conjugate X' of X . The proof of the latter was easy, but there does not seem to be such a straightforward proof of the new lemma. We have to prove that $\text{Stab}_1(p) = G^{00}$ for each μ -weakly random global type p . Passing to a bigger monster model M' , and arguing as in 5.8 of [33], μ has an extension to a definable left G^{00} -invariant measure μ' over M' . But then clearly there is a small M'_0 such that for all $g \in G(\bar{M}')$, $g\mu'$ does not fork over M'_0 . Now our μ -weakly random type p of μ extends to a μ' -weakly random type p' . By what we have just seen, p' is left f -generic (every left translate does not fork over a fixed M'_0). By Proposition 5.6(i) of [33], $\text{Stab}_1(p') = G^{00}(\bar{M}')$. It follows that $\text{Stab}_1(p) = G^{00}$. \square

Proposition 3.5.7. *Let G be a definable group. Then the following are equivalent :*

- (i) *There is a unique left G^{00} -invariant global Keisler measure of G lifting Haar measure on G/G^{00} ,*
- (ii) *Let \mathcal{P} be the family of global complete types of G such that $\text{Stab}_1(p) = G^{00}$. Let π be the canonical surjective map from \mathcal{P} to G/G^{00} , and h Haar measure on G/G^{00} . Then \mathcal{P} is nonempty and is dominated by $(G/G^{00}, \pi, h)$.*

Proof. (i) implies (ii) : This is like LHS implies RHS in 5.4. The nonemptiness of \mathcal{P} is given by the previous Lemma. Write C for G/G^{00} . For $c \in C$, let \mathcal{P}_c be those members of \mathcal{P} which concentrate on $\pi^{-1}(c)$. Suppose for a contradiction that for some definable subset X of G , the (closed) subset D of C consisting of c such that both $\mathcal{P}_c \cap X$ and $\mathcal{P}_c \cap X^c$ is non empty, has positive Haar measure. Then, as in proofs of 5.4 and 5.5, h lifts to two global left G^{00} -invariant Keisler measures, one giving $D \cap X$ positive measure (in fact that of D) and the other giving it 0 measure. Contradiction.

(ii) implies (i). The existence of μ is given by 5.5. So let us fix global left G^{00} -invariant μ which lifts Haar measure h on G/G^{00} . Let \mathcal{P}' be the set of family of μ -weakly random global types. By Lemma 5.6, \mathcal{P}' is a nonempty subset of \mathcal{P} which clearly maps onto $G/G^{00} = C$ under π . So now write π for the map $\mathcal{P}' \rightarrow C$. The assumption (ii) implies that also \mathcal{P}' is dominated by (C, π, h) . As in the proof of RHS implies LHS of Theorem

5.4, we conclude that for any definable set X , $\mu(X) = h(D)$ where $D = \{c \in C : \pi^{-1}(c) \subseteq X\}$. So μ is determined. \square

By 4.3, the previous proposition applies to any definable group with fsg. In fact the class \mathcal{P} of global types with stabilizer G^{00} can clearly be replaced by the subclass \mathcal{P}_{gen} of global *generic* types p of G . (Here generic is in the sense of [33].) Hence we have “generic compact domination” for fsg groups :

Proposition 3.5.8. *Suppose G has fsg. Let \mathcal{P}_{gen} be the space of global generic types of G , $\pi : G \rightarrow G/G^{00}$ as before and h Haar measure on G/G^{00} . Then \mathcal{P}_{gen} is dominated by $(G/G^{00}, \pi, h)$.*

Finally we point out that under an additional hypothesis on G , we can slightly strengthen the domination statement. Let us fix a definable group G , $\pi : G \rightarrow G/G^{00} = C$, and definable subset X of G . We will say that X is *left generic in $\pi^{-1}(c)$* if finitely many left translates of X by elements of G^{00} cover $\pi^{-1}(c)$.

We will be interested in the following hypothesis on an fsg group G :

(H) : Let $X \subseteq G$ be definable. Then X is generic in G if and only if for some small model M every left translate gX of X does not divide over M .

The left hand side implies the right hand side in any fsg group. We do not know an example of an fsg group G where (H) fails. It is true in any o-minimal expansion of RCF.

Lemma 3.5.9. *Suppose that the fsg group G satisfies (H). Let X be a definable subset of G , and $c \in C$. The following are equivalent :*

- (i) X is generic in $\pi^{-1}(c)$,
- (ii) $X \in p$ for every $p \in \mathcal{P}_{\text{gen}}$ concentrating on c ,
- (iii) for some definable set Y containing $\pi^{-1}(c)$, $Y \setminus X$ is not generic in G .

Proof. Without loss of generality let c be the identity of G/G^{00} .

(i) implies (ii) holds without even assuming (H) as the stabilizer of any generic type is G^{00} . (See [32] or [33].)

(ii) implies (iii) also holds without assuming (H) : By (ii), “ $x \in G^{00}$ ” \cup “ $x \notin X$ ” \cup “ $\neg\psi : \psi$ over \bar{M} , ψ nongeneric” is inconsistent, so by compactness, for some definable Y containing G^{00} , $Y \setminus X$ is nongeneric (we use here that the set of nongenerics is an ideal).

(iii) implies (i) : Let $Z = Y \setminus X$. Let M be a model over which X and Y are defined. By (H) there is some M -indiscernible sequence $(a_i : i < \omega)$ of elements of G such that $\bigcap_{i=1, \dots, n} a_i Z = \emptyset$ for some n . Let $g_i = a_1^{-1} a_i$ for $i = 1, \dots, n$. So $\bigcap_{i=1, \dots, n} g_i Z = \emptyset$. All elements of a_i are in the same coset of G^{00} , hence $g_i \in G^{00}$ for $i = 1, \dots, n$. As each $g_i Y$ contains G^{00} it follows that $g_1 X \cup \dots \cup g_n X \supseteq G^{00}$. \square

Corollary 3.5.10. *Suppose that G is a group with fsg which satisfies (H), and $\pi : G \rightarrow G/G^{00}$. Then for any definable subset X of G , for almost all $c \in G/G^{00}$ in the sense of Haar measure, either X is generic in $\pi^{-1}(c)$ or $\neg X$ is generic in $\pi^{-1}(c)$.*

Proof. Clear. □

§3.6 Borel measures over standard models

In this section we give a rich source of smooth measures in the case of theories of \mathbf{o} -minimal expansions of \mathbb{R} , as well as $\text{Th}(\mathbb{Q}_p)$. If M_0 is the standard model, $V \subseteq M_0^n$ is definable, and μ^* is a Borel probability measure on the topological space V , then by restricting μ^* to definable sets, we have a Keisler measure which we call μ , over M_0 . We will show that any such μ is smooth : has a unique extension to a Keisler measure μ' over a saturated model. It follows in particular that μ' will be “definable” ([32]), from which one can easily obtain “approximate definability” of μ^* in the sense of Karpinski and Macintyre [35], thereby considerably generalizing their results on approximate definability of the real and p -adic Haar measures on unit discs.

For now, T is an arbitrary complete theory. It is convenient to formally weaken the notion of a Keisler measure by allowing values in $[0, r]$ for some r , but of course maintaining finite additivity. Sometimes we may say that the Keisler measure μ is ON the definable set X if $\mu(X^c) = 0$.

Definition 3.6.1. Let μ_x be a Keisler measure over a model M . We will say that μ is *countably additive* over M , if whenever X is definable over M , Y_i are definable over M for $i < \omega$ and pairwise disjoint and $X(M)$ is the union of the $Y_i(M)$, then $\mu(X) = \sum_{i < \omega} \mu(Y_i)$.

Remark 3.6.2. (i) Of course the definition depends on M . When M is ω -saturated, any Keisler measure over M is countably additive, because if $X(M)$ is the union of the $Y_i(M)$ then by compactness it will be a finite subunion.

(ii) If, as above, M_0 is a structure whose underlying set is a topological space X say, and such that all definable subsets of the universe are Borel, THEN any Borel measure μ^* on X such that $\mu^*(X) \neq \infty$ induces a countably additive Keisler measure over M_0 (on X), by restricting to definable sets.

Theorem 3.6.3. *Let M_0 be either an \mathbf{o} -minimal expansion of $(\mathbb{R}, +, \cdot)$, or the structure $(\mathbb{Q}_p, +, \cdot)$. Let V be a definable set in M_0 , and μ a countably additive Keisler measure on V over M_0 . THEN μ is smooth. That is, μ has a unique extension to a global Keisler measure on V .*

Proof. We will distinguish here between the definable set V (as a functor say) and the set $V(M_0)$ of M_0 -points. The proof is by induction on the \mathbf{o} -minimal/ p -adic dimension of V (or $V(M_0)$) which we take to be n . Clearly it suffices to partition V into M_0 -definable sets V_1, \dots, V_k and prove the proposition for $\mu_i = \mu|_{V_i}$ for each i (where we stipulate that μ_i is 0 outside V_i). So by cell-decomposition we may assume that $V \subseteq I^n$ where I is the closed unit interval $[0, 1]$ in the \mathbf{o} -minimal case, and the valuation ring in the p -adic case. So in fact there is no harm in assuming that $V = I^n$.

Let us fix an extension μ' of μ to a global Keisler measure. And let D be a definable (over \bar{M}) subset of V . We aim to show that $\mu'(D)$ is determined, namely can be computed in terms of μ .

Recall that we have the standard part map st from $I^n(\bar{M})$ to $I^n(M_0)$, namely from $V(\bar{M})$ to $V(M_0)$. In both the \mathcal{o} -minimal and \mathbf{p} -adic cases all types over the standard model are definable ([40], [17]). As explained in [43] for example, this implies that for any definable in \bar{M} subset X of I^n , $\text{st}(X)$ is a definable set in the structure M_0 . In particular $\text{st}(D)$, $\text{st}(D^c)$, and the intersection $\text{st}(D) \cap \text{st}(D^c)$ are definable sets in the structure M_0 . Hence we can write $V(M_0)$ as the disjoint union of $Y(M_0)$, $D_0(M_0)$, and $D_1(M_0)$, where as the notation suggests Y, D_0, D_1 are definable over M_0 , $Y(M_0) = \text{st}(D) \cap \text{st}(D^c)$, $D_1(M_0) = \text{st}(D) \setminus Y$ and $D_0(M_0) = \text{st}(D^c) \setminus Y$. Note also that D_0, D_1 are open M_0 -definable subsets of V , and of course V is the disjoint union of Y, D_0 and D_1 .

Claim 1. The M_0 -definable subset $Y \cup (\text{cl}(D_0) \cap \text{cl}(D_1))$ of V has dimension $< n$.

Proof of Claim 1. Otherwise it contains an open M_0 definable set U say. But then either $D \cap U$ or $D^c \cap U$ contains an open M_0 -definable subset (of $V = I^n$). (In the \mathbf{p} -adic case this is Theorem 2.2(ii) of [43]. It is well-known in the \mathcal{o} -minimal case too, but formally follows from 10.3 of [32] for example.) But this is clearly impossible. For if, for example, W is an open M_0 -definable set contained in D , then for all $w \in W(M_0)$, $\text{st}^{-1}(w) \subseteq D$, so for each $w \in W(M_0)$, $w \notin Y(M_0)$, and $w \notin \text{cl}(D_0)(M_0)$. The claim is proved.

Let D_2 be the (closed, M_0 -definable) set $Y \cup (\text{cl}(D_0) \cap \text{cl}(D_1))$. Let μ_2 be $\mu|_{D_2}$. Namely μ_2 agrees with μ on M_0 -definable subsets of D_2 and is 0 on the complement of D_2 . Likewise define μ'_2 to be equal to μ' on definable subsets of D_2 and 0 on the complement of D_2 . Then as μ_2 is still countably additive, we see, by induction hypothesis and Claim 1, that μ'_2 is the unique global extension of μ_2 . In particular we have :

Claim 2. $\mu'_2(D) = \mu'(D \cap D_2)$ is determined.

Claim 3. Let $D_3(M_0)$ be an open M_0 -definable neighbourhood of the closed set $D_2(M_0)$. Then $D \setminus D_3 = D_1 \setminus D_3$, hence $\mu'(D \setminus D_3) = \mu(D_1 \setminus D_3)$.

Proof of Claim 3. Let $a \in V = V(\bar{M})$, and suppose $a \notin D_3$. So

(*) $\text{st}(a) \notin D_2(M_0)$.

Case (i) : $a \in D_1$. Then $\text{st}(a) \in \text{cl}(D_1)(M_0)$. By (*), $\text{st}(a) \notin Y$, and $\text{st}(a) \notin D_0$. Hence $\text{st}(a) \in D_1$ and we conclude that $a \in D$.

Case (ii) : $a \in D_0$. As in Case (i) we conclude that $\text{st}(a) \in D_0$ hence $a \notin D$.

This proves Claim 3.

Claim 4. $\mu'(D \setminus D_2) = \mu(D_1 \setminus D_2)$.

Proof. For small $\delta > 0$ let $D_\delta(M_0)$ be the δ neighbourhood of D_2 . Then $\cap_\delta D_\delta = D_2$. So $\mu(D_\delta \setminus D_2) \rightarrow 0$ as $\delta \rightarrow 0$, hence also $\mu'(D \cap (D_\delta \setminus D_2)) \rightarrow 0$ as $\delta \rightarrow 0$. It follows, using Claim 3, that $\mu'(D \setminus D_2) = \lim_{\delta \rightarrow 0} \mu'(D \setminus D_\delta) = \lim_{\delta \rightarrow 0} \mu(D_1 \setminus D_\delta) = \mu(D_1 \setminus D_2)$.

Claims 2 and 4 show that $\mu'(D)$ is determined. \square

Remark 3.6.4. (i) The key point about countably additive Keisler measures μ over the standard model is that any global extension μ' must assign 0 to definable sets which are

“infinitesimal”.

(ii) The inductive proof of Theorem 6.3 yields the following : for any definable subset D of $I^n(\bar{M})$, there is a partition of I^n into M_0 -definable cells V_1, \dots, V_k , such that for each i , EITHER for all $a \in V_i(M_0)$ $st^{-1}(a) \cap V_i \subseteq D$, OR for all $a \in V_i(M_0)$ $st^{-1}(a) \cap V_i \cap D = \emptyset$.

Constructions de mesures génériquement stables

§4.1 Introduction

A major theme in the study of NIP theories is the investigation of stable phenomena in them. In [55] Shelah defines *stable types* as types that are both finitely satisfiable and definable over a base. In [33], they are renamed *generically stable types* (and we will use the latter terminology) and a systematic study is made. This notion was generalized to Keisler measures in [29]. Keisler measures (introduced first in [37] and studied again with a different approach in [32], [33] and [29]) are regular Borel probability measures on the space of types. Equivalently, they are finitely additive probability measures on the boolean algebra of definable sets. They can also be considered as types in continuous logic (see [72]). Some basic facts about them are recalled in the first section.

After Keisler's seminal work [37], measures were introduced again for the study of NIP theories in Hrushovski, Peterzil and Pillay's paper [32], focussing mainly on invariant measures of groups. The notion of *fsg* group is introduced. Measures are studied more at depth in [33] in order to prove *compact domination* for definably compact groups in o-minimal theories. Again, applications use only invariant measures on groups, but a general study is initiated. Also, the notion of generically stable type is defined and the question is asked of a generalization to measures. This is done in [29] by Hrushovski, Pillay and the author where generically stable measures are defined. A number of equivalent properties are given.

The following examples of generically stable measures are known :

- a generically stable type : this is the motivating example,
- the (translation) invariant measure of an *fsg* group,
- the A -invariant measure extending an *fsg* type over A ([29] Section 4),
- the average measure of an indiscernible segment ([29], Proposition 3.7),
- the Keisler measure induced by a σ -additive measure on the standard model, \mathbb{R} or \mathbb{Q}_p ([29], Section 6). In fact, those measures are proven to be *smooth*, a stronger property.

In this paper, we generalize the last two constructions in this list. First, we show how to *symmetrize* any measure, or equivalently to average an indiscernible segment of measures. Second, we show that any σ -additive measure induces a generically stable measure, under the assumption that externally definable sets in two variables are measurable (so that Fubini applies). Smoothness of the induced measure is not true in general, and to recover the full result of [29] we also include a proof that all generically stable measures in \mathbb{R} or \mathbb{Q}_p are smooth. In order to prove those results, we first establish in Section 2 a criterion to recognize a product measure $\mu_x \otimes \lambda_y$ when μ_x happens to be finitely satisfiable over some small model. Our strategy for proving that a given measure μ is generically stable will then be to construct a symmetric measure $\eta_{x_1 x_2 \dots}$ in ω variables, and show using this criterion that it is the Morley sequence of μ .

We will use standard notation. We will work with a complete first order theory T in some language L ; T is assumed to be NIP throughout the paper. For simplicity, we assume that T is one-sorted and work in T^{eq} . We have a monster model \bar{M} ; M, N, \dots will denote small submodels of \bar{M} , and A, B, C, \dots small parameter sets. We will not distinguish between points and tuples; they will be named by a, b, c, \dots and x, y, z, \dots will designate variables of finite or infinite tuples. The notation $L_x(A)$ denotes the set of formulas with parameters in A and free variable x .

The space of types over A in variable x is designated by $S_x(A)$. It is equipped with the usual compact topology and the associated σ -algebra of Borel subsets. By “ X is Borel over A ”, we mean that it is a Borel subset of some $S_x(A)$. We write $a \equiv_M b$ for $\text{tp}(a/M) = \text{tp}(b/M)$.

By a *global* type or measure, we mean a type or a measure over \bar{M} .

§4.2 Preliminaries

We recall some basic facts about Keisler measures. We will be brief, and the reader is referred to [33] and [29] for more details.

We make a blanket assumption that T is NIP.

Basic definitions

A *Keisler measure* (or simply a *measure*) over A in variable x is a finitely additive probability measure on the boolean algebra $L_x(A)$ of formulas with free variable x and parameters in A . As in section 4 of [33], such a measure extends uniquely to a regular Borel probability (σ -additive) measure on the type space $S_x(A)$. (“Regular” means that the measure of any Borel set X is the infimum of the measures of open sets O such that $X \subseteq O$. Furthermore, in our situation, working on a totally disconnected space, the measure of O is itself the supremum of measures of clopen sets inside it.)

Conversely, given a regular Borel measure on $S_x(A)$, its restriction to the clopen sets gives a Keisler measure.

We will denote by $\mathcal{M}_x(A)$ the space of Keisler measures over A and often write $\mu_x \in \mathcal{M}(A)$ for $\mu \in \mathcal{M}_x(A)$, keeping track of the variable in the name of the measure. We can consider $\mathcal{M}_x(A)$ as a subset of $[0, 1]^{L_x(A)}$. It inherits the product topology making it a compact Hausdorff space.

Lemma 4.2.1. *Let $\Omega \subseteq L_x(A)$ be a set of formulas closed under intersection, union and complement and containing \top . Let μ_0 be a finitely additive measure on Ω with values in $[0, 1]$ such that $\mu_0(\top) = 1$. Then μ extends to a Keisler measure over A .*

Proof. By compactness in the space $[0, 1]^{L_x(A)}$, it is enough to show that given formulas $\psi_1(x), \dots, \psi_n(x)$ in $L_x(A)$, there is a function $f : \langle \psi_1, \dots, \psi_n \rangle \rightarrow [0, 1]$ finitely additive and compatible with μ_0 (where $\langle B \rangle$ denotes the boolean algebra generated by B). We may assume that ψ_1, \dots, ψ_n are the atoms of the boolean algebra B that they generate.

The elements of Ω in B form a sub-boolean algebra. Let ϕ_1, \dots, ϕ_m be its atoms. We have say :

$$\phi_1 = \psi_{i_1(1)} \vee \dots \vee \psi_{i_1(l_1)} \quad \phi_2 = \psi_{i_2(1)} \vee \dots \vee \psi_{i_2(l_2)} \quad \text{etc.}$$

Then any finitely additive f satisfying $f(\psi_{i_1(1)}) + \dots + f(\psi_{i_1(l_1)}) = \mu_0(\phi_1)$ etc. will do. \square

Here are some basic definitions :

Definition 4.2.2. Let $M \prec N$, with N $|M|^+$ -saturated and let $\mu_x \in \mathcal{M}(N)$,

- μ is *finitely satisfiable* in M if for every $\phi \in L_x(N)$ such that $\mu(\phi) > 0$, there is $a \in M$ such that $N \models \phi(a)$.
- μ is *M-invariant* if for every $\phi(x; y) \in L$, and $b \equiv_M b'$, $\mu(\phi(x; b)) = \mu(\phi(x; b'))$.
- μ is *definable* over M if it is M -invariant and for every $\phi(x; y) \in L$, and $r \in [0, 1]$, the set $\{p \in S_y(M) : \mu(\phi(x; b)) \leq r \text{ for any } b \in N, b \models p\}$ is a closed set of $S_y(M)$.
- μ is *Borel-definable* over M if the above set is a Borel set of $S_y(M)$.

Proposition 4.2.3 ([33] 4.9). *If $\mu \in \mathcal{M}(N)$ is M -invariant (N is $|M|^+$ -saturated), then it is Borel definable over M .*

The *support* of $\mu \in \mathcal{M}(N)$ is the set of types $p \in S(N)$ satisfying $p \vdash \neg \phi(x)$ for every $\phi(x) \in L(N)$ such that $\mu(\phi(x)) = 0$. We will denote the support of μ by $S(\mu)$; it is a closed set of $S(N)$. Note that if μ is finitely satisfiable in M then every type in $S(\mu)$ is also finitely satisfiable in M .

The next proposition says that in NIP theories, measures can be well approximated by averages of types. We use the notation $\text{Av}(T_i : i = 1 \dots n)$ which stands for $\frac{1}{n} \{i : T_i \text{ holds}\}$.

Proposition 4.2.4 ([29] 2.8). *Let μ_x be any measure over M . Let $\phi(x, y) \in L$, $\epsilon > 0$, and let X_1, \dots, X_k be Borel sets of $S_x(M)$. Then for all large enough n there are $p_1, \dots, p_n \in S_x(\mu)$ such that for all $r = 1, \dots, k$ and all $b \in M$, $\mu(X_r \cap \phi(x, b))$ is within ϵ of $\text{Av}(p_i \in X_r \text{ and } p_i \vdash \phi(x, b) : i = 1 \dots n)$.*

Note in particular that if μ_x is finitely satisfiable in some model N , then the types p_i 's are also finitely satisfiable in N .

Invariant extensions

As in the case of types, the study of Keisler measures differs from measure theory in that the space in two dimensions is not the product of the one dimensional spaces, and there are in general different ways to amalgamate two measures in one variable to form a measure in two variables. We recall here the basic construction of invariant extensions.

Let $M \prec N$, N being $|M|^+$ -saturated, and $\mu_x \in \mathcal{M}(N)$ be M -invariant. If $\lambda_y \in \mathcal{M}(N)$ is any measure, then we can define the *invariant extension* of μ_x over λ_y , denoted $\mu_x \otimes \lambda_y$. It is a measure in the two variables x, y defined the following way. Let $\phi(x, y) \in L(N)$. Take a small model $P \prec N$ containing M and the parameters of ϕ . Define $\mu_x \otimes \lambda_y(\phi(x, y)) = \int f(p) d\lambda_y$, the integral ranging over $S_y(P)$ where $f(p) = \mu_x(\phi(x, b))$ for $b \in N$, $b \models p$ (this function is Borel-measurable by Borel-definability of μ_x). It is easy to check that this does not depend on the choice of P .

If λ_y is also invariant, we can also form the product $\lambda_y \otimes \mu_x$. In general it will not be the case that $\lambda_y \otimes \mu_x = \mu_x \otimes \lambda_y$.

If μ_x is a global M -invariant measure, we define by induction : $\mu_{x_1 \dots x_n}^{(n)}$ by $\mu_{x_1}^{(1)} = \mu_{x_1}$ and $\mu_{x_1 \dots x_{n+1}}^{n+1} = \mu_{x_{n+1}} \otimes \mu_{x_1 \dots x_n}^{(n)}$. We let $\mu_{x_1 x_2 \dots}^{(\omega)}$ be the union and call it the *Morley sequence* of μ_x . It is an indiscernible sequence in the following sense.

Definition 4.2.5. A measure $\mu_{x_1 x_2 \dots}$ is indiscernible over A if for every $\phi(x_1, \dots, x_n) \in L(A)$ and indices $i_1 < \dots < i_n$, we have

$$\mu(\phi(x_1, \dots, x_n)) = \mu(\phi(x_{i_1}, \dots, x_{i_n})).$$

We define in the same way $\mu_{x_1 x_2 \dots}$ to be *totally indiscernible* by removing in the above definition the assumption that the indices i_1, \dots, i_n are ordered.

Similarly, given I any linear order, we can define $\mu_{\bar{x}}^{(I)}$, where $\bar{x} = \langle x_t : t \in I \rangle$, the Morley sequence of μ_x indexed by I .

We will need the following result from [71] (see also [29], 2.10).

Proposition 4.2.6. *If $\mu_{x_1, x_2, \dots} \in \mathcal{M}(A)$ is indiscernible (over \emptyset), then for every formula $\phi(x, y) \in L$ and $b \in M$, $\lim_{i \rightarrow \omega} \mu(\phi(x_i, b))$ exists.*

Equivalently, for any $\phi(x, y)$ and $\epsilon > 0$, there is N such that for any indiscernible $\mu_{x_1, x_2, \dots}$, we have $|\mu(\phi(x_i, y)) - \mu(\phi(x_{i+1}, y))| \geq \epsilon$ for at most N values of i .

Generically stable measures

This paper is concerned with building generically stable measures. They are measures that behave very much like types in a stable theory (at least as far as non-forking extensions are concerned). Generically stable *types* are studied in [33] (and previously by Shelah in [55]) and this notion was naturally extended to measures in [29]. We recall here some equivalent definitions (a few others are given in [29], Theorem 3.3).

Theorem 4.2.7 (Generically stable measure). *Let μ_x be a global M -invariant measure. Then the following are equivalent :*

1. μ_x is both definable and finitely satisfiable (necessarily over M),
2. $\mu_{x_1, x_2, \dots}^{(\omega)}|_M$ is totally indiscernible,
3. for any global M -invariant Keisler measure λ_y , $\mu_x \otimes \lambda_y = \lambda_y \otimes \mu_x$,
4. μ commutes with itself : $\mu_x \otimes \mu_y = \mu_y \otimes \mu_x$.

If μ_x satisfies one of those properties, we say it is generically stable.

Let μ_x be a measure, and f a definable map whose domain is in the sort of x . Then one can consider the push-forward measure $\lambda_y = f_*(\mu_x)$ defined by $\lambda_y(\phi(y)) = \mu_x(\phi(f(x)))$. This is again a Keisler measure. If μ_x is definable over M (resp. finitely satisfiable in M) and f is M -definable, then $f_*(\mu_x)$ is again definable over M (resp. finitely satisfiable in M). In particular, if μ_x is a global generically stable measure, then $f_*(\mu_x)$ is also generically stable.

Smooth measures

Definition 4.2.8 (Smooth). A measure $\mu \in \mathcal{M}(N)$ is smooth if μ has a unique global extension. If $M \subset N$, we will say that μ is smooth over M if $\mu|_M$ is smooth.

The following important properties are proved in [29].

Proposition 4.2.9 ([29], 2.3). Let μ_x be smooth over M and let $\phi(x, y) \in L$ and $\epsilon > 0$. Then there are formulas $\nu_i^1(x)$, $\nu_i^2(x)$ and $\psi_i(y)$ for $i = 1, \dots, n$ in $L(M)$ such that :

1. The formulas $\psi_i(y)$ form a partition of the y -space,
2. for all i and $b \in \bar{M}$, if $\psi_i(b)$ holds, then $\bar{M} \models \nu_i^1(x) \rightarrow \phi(x, b) \rightarrow \nu_i^2(x)$,
3. for each i , $\mu_x(\nu_i^2(x)) - \mu_x(\nu_i^1(x)) < \epsilon$.

Note that conversely, if the conclusion holds for all $\phi(x, y)$ and ϵ , then μ is smooth.

Corollary 4.2.10. If μ_x is smooth over N , then :

1. there is $M \prec N$ of size $|T|$ such that μ_x is smooth over M ,
2. μ_x is definable and finitely satisfiable in N (in particular μ_x is generically stable),
3. if λ_y is a measure over N , then there is a unique separated amalgam ω_{xy} of μ_x and λ_y (see Definition 4.3.1).

Will also need the following fact (initially from [37]), that we consider as a way to realize measures.

Lemma 4.2.11 ([29], 2.2). Let μ_x be a measure over M . Then there is an extension $M \prec N$ and a measure μ'_x over N extending μ_x such that μ'_x is smooth.

Some additional facts about smooth measures can be found in Section 4.5.

Two basic examples

We describe the canonical example of a generically stable measure we have in mind. Consider the following theory : the signature is $\{\leq, E\}$. The reduct to $\{\leq\}$ is a dense linear order, and the theory says that E is an equivalence relation with infinitely many classes, each of which is dense with respect to \leq . This theory has no generically stable type over the main sort (because every type falls in some cut of the linear order). However, one can build a generically stable measure over the main sort by averaging types that fall in different cuts. More precisely, assume for example that we work over a model (M, \leq, E) and we have an increasing embedding $f : ((0, 1), \leq) \rightarrow (M, \leq)$, where $(0, 1)$ denotes the standard unit open interval. Let λ_0 denote the standard Lebesgue measure on $(0, 1)$. Define a Keisler measure μ on M by $\mu(a \leq x) = \lambda_0(f^{-1}([a, +\infty)))$ and $\mu(aEx) = 0$ for all $a \in M$. This measure is generically stable. Let π be the canonical projection from M to M/E , then the push-forward $\pi_*(\mu)$ is a type, namely the unique non-realized type p of M/E (we say that μ lifts p). We will see that this phenomenon is general : we can always lift generically stable types in an imaginary sort to generically stable measures on the main sort (see Lemma 4.5.4). Note that in this example, it is not possible to lift p to a generically stable *type*.

Of course this example is rather special in that the generic type of M/E is stable and not just generically stable. Here is an example where this is not the case. Start with the structure with universe \mathbb{Q} and with language $\{P_n(x, y) : n < \omega\}$ where $P_n(x, y)$ holds if and only if $x < y \wedge |x - y| < n$. Call M_0 this structure. Then the set of formulas $\{\neg P_n(x, a) : a \in M_0\}$ defines a complete type p over M_0 and this type is generically stable. Now expand that structure with a linear order $<$ such that every infinite definable set of M_0 is dense-co-dense with respect to $<$. Let M be the resulting structure. Then the type p extends in many ways to a generically stable measure μ on M : the reduct of μ to $<$ can be any Lebesgue measure as above. Such a μ is a non-smooth generically stable measure, but there are no generically stable types in the expanded theory, even in M^{eq} .

§4.3 Amalgams

We fix throughout this section the following objects : $M \prec N$ two models, with N being $|M|^+$ -saturated, and μ_x, λ_y two measures over N . An *amalgam* of μ_x and λ_y is a measure ω_{xy} extending $\mu_x \cup \lambda_y$. We are interested in the characterization of different possible amalgams, especially ‘free’ amalgams.

The most basic property an amalgam can have is independence in the sense of probability theory, which we call separation.

Definition 4.3.1 (Separated). The amalgam ω_{xy} is separated if

$$\omega_{xy}(\phi(x) \wedge \psi(y)) = \mu_x(\phi(x)) \cdot \lambda_y(\psi(y))$$

for all $\phi(x), \psi(y) \in L(N)$.

Note that if μ_x is M -invariant, then $\mu_x \otimes \lambda_y$ is separated. Also if μ_x or λ_y is a type, then every amalgam is separated.

We now go on to define when the amalgam ω_{xy} is a finitely satisfiable extension of μ_x . A natural attempt would be to ask for example that $\omega_{xy}(\theta(x, y)) \leq \sup_{a \in M} \lambda_y(\theta(a, y))$. However, this seems to be too weak, and we will ask for something stronger, allowing to ‘localize’ on any clopen $\phi(x)$.

Definition 4.3.2 (fs extension). The amalgam ω_{xy} is a fs extension in M of μ_x over λ_y if the following holds for all $\theta(x, y), \phi(x) \in L(N)$:

$$\omega_{xy}(\theta(x, y) \wedge \phi(x)) \leq \mu_x(\phi(x)) \cdot \sup_{a \in \phi(M)} \lambda_y(\theta(a, y)).$$

First some basic observations.

Fact 4.3.3. 1. The existence of a fs extension in M of μ_x over λ_y implies that μ_x itself is finitely satisfiable in M (hence M -invariant).

2. If ω_{xy} is a fs extension in M of μ_x , then it is a separated amalgam.

3. Assume λ_y is a type realized by some b . We can view ω_{xy} as $\omega'_x \in \mathcal{M}(Nb)$. Then ω_{xy} is a fs extension in M of μ_x if and only if ω'_x is finitely satisfiable in M .

Proof. 1. Assume there is such an amalgam ω_{xy} . Let $\phi(x) \in L_x(N)$ be such that $\phi(M) = \emptyset$. Then applying Definition 4.3.2, with $\theta(x, y) = “x = x”$, we obtain

$$\mu_x(\phi(x)) = \omega_{xy}(\phi(x)) \leq \mu_x(\phi(x)) \cdot \sup_{a \in \phi(M)} \lambda_y(\theta(a, y)) = 0.$$

2. Apply the definition with $\theta(x, y) = \psi(y)$ and then $\theta(x, y) = \neg\psi(y)$.

3. Assume ω_{xy} is a fs extension of μ_x . Let $\theta(x, y) \in L_{xy}(N)$ be such that $\theta(M, b) = \emptyset$. Then

$$\omega'_x(\theta(x, b)) = \omega_{xy}(\theta(x, y)) \leq \mu_x(\phi(x)) \cdot \sup_{a \in \phi(M)} \lambda_y(\theta(a, y)) = 0.$$

This shows that ω'_x is finitely satisfiable in M .

Conversely, assume that ω'_x is finitely satisfiable in M and take $\theta(x, y), \phi(x) \in L(N)$. Assume $\omega_{xy}(\theta(x, y) \wedge \phi(x)) > 0$. Then by hypothesis, there is $a \in M$ such that $a \models \theta(x, b) \wedge \phi(x)$. For that a , we have $\lambda_y(\theta(a, y)) = 1$. Therefore :

$$\omega_{xy}(\theta(x, y) \wedge \phi(x)) \leq \omega_{xy}(\phi(x)) = \mu_x(\phi(x)) \cdot \sup_{a \in \phi(M)} \lambda_y(\theta(a, y)).$$

And ω_{xy} is a fs extension of μ_x over λ_y . □

As before, we use the notation $\text{Av}(T_i : i = 1, \dots, n)$ to mean $\frac{1}{n} |\{i : T_i \text{ holds}\}|$.

Proposition 4.3.4. Assume that μ_x is finitely satisfiable in M , then $\omega_{xy} = \mu_x \otimes \lambda_y$ is a fs extension (in M) of μ_x over λ_y .

Proof. We first assume that $\mu = p$ is a type. Let $\theta(x, y) \in L(N)$ and $\phi(x) \in L(N)$ such that $p \vdash \phi(x)$, and take a small model $P \subset N$ containing M and the parameters of θ .

Then $\omega(\theta(x, y)) = \lambda(B)$ where $B \subseteq S(P)$ is the Borel subset $B = \{q : p \vdash \theta(x, b) \text{ for some (any) } b \models q\}$. Let $\epsilon > 0$. By Proposition 4.2.4, there are points $b_1, \dots, b_n \in N$ such that :

$$|\lambda(B) - \text{Av}(\text{tp}(b_i/P) \in B : i = 1..n)| \leq \epsilon,$$

$$\forall a \in P, |\lambda(\theta(a, y)) - \text{Av}(\theta(a, b_i) : i = 1..n)| \leq \epsilon.$$

As p is finitely satisfiable in M , there is $a_0 \in \phi(M)$ such that for every $i \in \{1, \dots, n\}$:

$$p \vdash \theta(x, b_i) \leftrightarrow N \models \theta(a_0, b_i).$$

Now $p \vdash \theta(x, b_i) \iff \text{tp}(b_i/P) \in B$ so :

$$\begin{aligned} \lambda(B) &\approx \text{Av}(\text{tp}(b_i/P) \in B : i = 1..n) \\ &= \text{Av}(p \vdash \theta(x, b_i) : i = 1..n) \\ &= \text{Av}(\theta(a_0, b_i) : i = 1..n) \\ &\approx \lambda(\theta(a_0, y)). \end{aligned}$$

(Where $x \approx y$ means $|x - y| \leq \epsilon$.)

So $|\lambda(B) - \lambda(\theta(a_0, y))| \leq 2\epsilon$. As this is true for all $\epsilon > 0$, and remembering $\lambda(B) = \omega(\theta(x, y))$, we deduce $\omega(\theta(x, y)) \leq \sup_{a \in \phi(M)} \lambda(\theta(a, y))$. This finishes the proof in the case $\mu = p$.

We now consider the general case. Let as above $\theta(x, y), \phi(x) \in L(P)$. Let $\epsilon > 0$. By Proposition 4.2.4 and the remark after it, we can find $p_1, \dots, p_n \in S_x(N)$ finitely satisfiable in M and such that :

$$\forall b \in N, |\mu(\theta(x, b)) - \text{Av}(p_i \vdash \theta(x, b))| \leq \epsilon$$

and

$$|\mu(\phi(x)) - \text{Av}(p_i \vdash \phi(x))| \leq \epsilon.$$

Let, for $b \in M$, $f(b) = \mu(\theta(x, b) \wedge \phi(x))$ and $f_n(b) = \frac{1}{n} \text{Card}\{k : p_k \vdash \theta(x, b) \wedge \phi(x)\}$. Let m be the number of indices i for which $p_i \vdash \phi(x)$.

Then,

$$\begin{aligned} \omega(\theta(x, y) \wedge \phi(x)) &= \int f(y) d\lambda \\ &\leq \int f_n(y) d\lambda + \epsilon \\ &\leq \frac{m}{n} \sup_k \lambda(\{b \mid p_k \vdash \theta(x, b) \wedge \phi(x)\}) + \epsilon \\ &\leq \frac{m}{n} \sup_{a \in \phi(M)} \lambda(\theta(a, y)) + \epsilon \\ &\leq \mu(\phi(x)) \sup_{a \in \phi(M)} \lambda(\theta(a, y)) + 2\epsilon. \end{aligned}$$

(We use the first part of the proof to go from line 3 to 4).

As ϵ was arbitrary, we are done. \square

We now show that if μ_x is finitely satisfiable in M , then the invariant extension is the only fs extension of μ_x over λ_y .

Proposition 4.3.5. *Assume that μ_x is finitely satisfiable in M and ω_{xy} is a fs extension in M of μ_x , then $\omega_{xy} = \mu_x \otimes \lambda_y$.*

Proof. Let $\theta(x, y) \in L(N)$ and let $\epsilon > 0$. Let $P \subset N$ be a small model containing M and the parameters of θ . By Proposition 4.2.4 we can find $b_1, \dots, b_n \in N$ such that

$$|\lambda(\theta(a, y)) - \text{Av}(\theta(a, b_i))| \leq \epsilon, \text{ for all } a \in P.$$

For all $k \in \{0, \dots, n\}$, let $B_k(x)$ be the formula saying : “there are exactly k values of i for which $\theta(x, b_i)$ is true”.

Then :

$$\begin{aligned} \omega(\theta(x, y)) &= \sum_k \omega(\theta(x, y) \wedge B_k(x)) \\ &\leq \sum_k \mu(B_k(x)) \cdot \sup_{a \in B_k(M)} \lambda(\theta(a, y)) \\ &\leq \sum_k \frac{k}{n} \mu(B_k(x)) + \epsilon. \end{aligned}$$

(We use finite satisfiability of μ on line 2).

Similarly,

$$\omega(\neg\theta(x, y)) \leq \sum_k \left(1 - \frac{k}{n}\right) \mu(B_k(x)) + \epsilon,$$

and therefore

$$\left| \omega(\theta(x, y)) - \sum_k \frac{k}{n} \mu(B_k(x)) \right| \leq \epsilon.$$

Letting $\epsilon \rightarrow 0$, we see that $\omega(\theta(x, y))$ is uniquely determined by μ and λ and the fact that ω is a fs extension of μ . By the previous proposition, we have $\omega_{xy}(\theta(x, y)) = \mu_x \otimes \lambda_y(\theta(x, y))$. \square

§4.4 Symmetrizations

The following construction already appears in [29]. Let $I = \langle a_t : t \in [0, 1] \rangle$ be an indiscernible sequence indexed by $[0, 1]$ (we will call this an *indiscernible segment*). If $\phi(x, y)$ is a formula and $b \in M$, then NIP implies that $\lfloor \phi(x, b) \rfloor := \{t \in [0, 1] : \models \phi(a_t, b)\}$ is a finite union of intervals and points. Let m denote the Lebesgue measure on $[0, 1]$. Then we can define a global measure $\mu = \text{Av}(I)$ by $\mu(\phi(x, b)) = m(\lfloor \phi(x, b) \rfloor)$. It is called the average measure of I .

Lemma 4.4.1. *For any indiscernible segment $I \subset M$, the average measure $\mu = \text{Av}(I)$ is generically stable over M . Furthermore, for any n and $\phi(x_1, \dots, x_n) \in L(\bar{M})$,*

$$\mu^{(n)}(\theta(x_1, \dots, x_n)) = \int_{t_1 \in [0,1]} \dots \int_{t_n \in [0,1]} \theta(x_{a_1}, \dots, x_{a_n}) dt_1 \dots dt_n.$$

Proof. First notice that μ is finitely satisfiable in M by construction. It is easy to check directly from the definition that the formula given for $\mu^{(n)}$ is valid. From this, it follows that $\mu^{(n)}$ is symmetric for all n , therefore μ is generically stable. \square

If $p \in S(M)$ is any type. Let \tilde{p} be an M -invariant global extension of p (for example, a coheir). Let I be a Morley segment of \tilde{p} (i.e., a Morley sequence indexed by $[0, 1]$). Then $\text{Av}(I)$ is a generically stable measure and extends p . Note that if p is already generically stable, then $\mu = p$. In general, by the previous lemma, the Morley sequence of μ is a *symmetrization* of the Morley sequence of p , i.e. $\mu^{(n)}|_M$ is the average over all permutation of variables of $p^{(n)}|_M$:

$$\mu^{(n)}(\theta(x_1, \dots, x_n)) = \frac{1}{n!} \sum_{\sigma \in \mathfrak{S}_n} p^{(n)}(\theta(x_{\sigma(1)}, \dots, x_{\sigma(n)})),$$

for $\theta(x_1, \dots, x_n) \in L(M)$. For this reason, we will call μ a symmetrization of p .

Our aim now is to define the same construction starting with a measure instead of a type p . We start with $\mu \in \mathcal{M}_x(\mathfrak{C})$ a global measure invariant over some small model M . We consider its *Morley segment* $\mu_{\bar{x}}^{([0,1])}$ where \bar{x} stands for $\langle x_t : t \in [0, 1] \rangle$ ($\mu^{([0,1])}$ is the Morley sequence of μ indexed by $[0, 1]$ with the usual order). Now let $\nu_{\bar{x}}$ be a smooth extension of $\mu_{\bar{x}}^{([0,1])}|_M$ (this is the analogue of taking a realization of a Morley segment of p). For any $b \in M$, consider the function $f_{\phi(x,b)} : t \mapsto \nu_{\bar{x}}(\phi(x_t, b))$. By Proposition 4.2.6, this function has only countably many points of discontinuity. It is therefore integrable on $[0, 1]$.

Definition 4.4.2 (Symmetrization). Let $\mu, \mu_\Sigma \in \mathcal{M}_x(\mathfrak{C})$ be two measures and M a small model. We say that μ_Σ is a symmetrization of μ over M if there exists a measure $\nu \in \mathcal{M}_{\bar{x}}(\mathfrak{C})$ where $\bar{x} = \langle x_t : t \in [0, 1] \rangle$ such that :

- μ is invariant over M ,
- ν is a smooth extension of $\mu^{([0,1])}|_M$,
- for every $\theta(x; y) \in L$ and $b \in \mathfrak{C}$, we have :

$$\mu_\Sigma(\phi(x; b)) = \int_{t \in [0,1]} \nu(\phi(x_t, b)) dt.$$

In this case, we will say that μ_Σ is the symmetrization of μ over M built from ν .

Notice that the restrictions of μ_Σ and μ to the model M coincide.

Proposition 4.4.3. *If $\mu_\Sigma \in \mathcal{M}_\chi(\bar{M})$ is a symmetrization of a measure μ over M , then μ_Σ is generically stable. Furthermore, for every n , and $\theta(x_1, \dots, x_n) \in L(M)$,*

$$\mu_\Sigma^{(n)}(\theta(x_1, \dots, x_n)) = \frac{1}{n!} \sum_{\sigma \in \mathfrak{S}_n} \mu^{(n)}(\theta(x_{\sigma(1)}, \dots, x_{\sigma(n)})).$$

Proof. Assume that μ_Σ is a symmetrization of μ built from $\nu \in \mathcal{M}_{\bar{\chi}}(\mathfrak{C})$.

Let P be a small model containing M over which ν is smooth. Then ν is finitely satisfiable in P , and it follows that μ_Σ is also finitely satisfiable in P . For $n < \omega$, define an n -ary global measure $\lambda_{x_1 \dots x_n}^n$ by

$$\lambda^n(\theta(x_1, \dots, x_n)) = \int_{t_1 \in [0,1]} \dots \int_{t_n \in [0,1]} \nu(\theta(x_{t_1}, \dots, x_{t_n})) dt_1 \dots dt_n,$$

for any formula $\theta(x_1, \dots, x_n) \in L(\bar{M})$.

We will show by induction that $\lambda^n = \mu_\Sigma^{(n)}$ for all n . The second assertion of the proposition will follow immediately by direct computation (remembering $\mu^{(\omega)}|_M = \nu|_M$). Then the first assertion follows since the expression given for $\mu_\Sigma^{(n)}$ is symmetric. We show that λ^n defines a fs extension in P of μ_Σ , and is therefore, by Proposition 4.3.5 and induction equal to $\mu_\Sigma^{(n)}$.

For simplicity of notations, we write the details only for $n = 2$ (the case $n = 1$ being true by definition). Let $\theta(x, y), \phi(x) \in L(\bar{M})$. The transformations are explained below.

$$\begin{aligned} \lambda_{x,y}^2(\theta(x, y) \wedge \phi(x)) &= \int_{t \in [0,1]} \int_{t' \in [0,1]} \nu_{\bar{x}}(\theta(x_t, x_{t'}) \wedge \phi(x_t)) dt' dt \\ &= \int_{t \in [0,1]} \mu_{\Sigma,y} \otimes \nu_{\bar{x}}(\theta(x_t, y) \wedge \phi(x_t)) dt \\ &\leq \sup_{a \in \phi(P)} \mu_{\Sigma,y}(\theta(a, y)) \times \int_{t \in [0,1]} \nu_{\bar{x}}(\phi(x_t)) dt \\ &\leq \mu_{\Sigma,x}(\phi(x)) \cdot \sup_{a \in \phi(P)} \mu_{\Sigma,y}(\theta(a, y)). \end{aligned}$$

Explanation : $\nu_{\bar{x}}$ is a smooth measure, so by Corollary 4.2.10, it admits a unique separated amalgam with any other measure. In particular with $\mu_{\Sigma,y}$. The measure $\omega_{\bar{x}y}$ defined by $\omega_{\bar{x}y}(\theta(\bar{x}, y)) = \int_{t' \in [0,1]} \nu_{\bar{x}}(\theta(\bar{x}, x_{t'})) dt'$ is such an amalgam. Therefore it is equal to $\mu_{\Sigma,y} \otimes \nu_{\bar{x}}$. This justifies the second line.

As $\mu_{\Sigma,y}$ is finitely satisfiable in P , $\mu_{\Sigma,y} \otimes \nu_x$ is a fs extension of $\mu_{\Sigma,y}$ over ν_x (in P); this explains the third line. The forth line is just the definition of μ_Σ . □

Proposition 4.4.4. *Let $\mu \in \mathcal{M}_\chi(\bar{M})$ be a global M -invariant measure, and μ_Σ a symmetrization of μ over M . Let also f be an M -definable function whose domain and range are in the same sort as the variable x .*

1. If μ is generically stable, then $\mu_\Sigma = \mu$,
2. $f_*(\mu_\Sigma)$ is a symmetrization of $f_*(\mu)$ over M ,
3. If $f_*(\mu) = \mu$, then $f_*(\mu_\Sigma) = \mu_\Sigma$.

Proof. 1. If μ is generically stable, then $\mu^{(\omega)}$ is totally indiscernible. It follows by Proposition 4.2.6 that for every $\phi(x, c) \in L(\bar{M})$ and $\epsilon > 0$, the set $\{t \in [0, 1] : |\nu(\phi(a_t, c)) - \mu(\phi(a, c))| > \epsilon\}$ is finite. Therefore the definition of μ_Σ implies that $\mu_\Sigma = \mu$.

2. Clear : $f_*(\mu_\Sigma)$ is a symmetrization of $f_*(\mu)$ built using $f_*(\nu)$.

3. Let $\nu \in S_{\bar{x}}(\mathcal{C})$ be such that μ_Σ is a symmetrization of μ built from ν . For $I \subseteq [0, 1]$, define the measure $\nu^I \in S_{\bar{x}}(\mathcal{C})$ by $\nu^I(\phi(x_{t_1}, \dots, x_{t_n})) = \nu(\phi(f_{t_1}^I(x_{t_1}), \dots, f_{t_n}^I(x_{t_n})))$ where $\phi(x_1, \dots, x_n) \in L_{\bar{x}}(\bar{M})$ and $f_t^I = f$ if $t \in I$ and is the identity otherwise.

Claim 1 : For every I , ν^I is indiscernible.

Proof. The claim concerns only the restriction to M (indeed to \emptyset) of the measure ν^I . Note that $\nu|_M$ is just $\mu^{([0,1])}|_M$, and as f is M definable, the property we need to check does not depend on the choice of the smooth extension ν . Let $\theta(x_1, \dots, x_n) \in L(M)$ and $i_1, \dots, i_n \in [0, 1]$. We want to see that $\mu^{([0,1])}(\theta(x_{i_1}, \dots, x_{i_n})) = \mu^{([0,1])}(\theta(f^I(x_{i_1}), \dots, f^I(x_{i_n})))$. To simplify notations, we note that $\mu^{([0,1])}(\theta(x_{i_1}, \dots, x_{i_n})) = \mu^{(n)}(\theta(x_1, \dots, x_n))$. So it is enough to show that $\mu^{(n)}(\theta(x_1, \dots, x_n)) = \mu^{(n)}(\theta(g_1(x_1), \dots, g_n(x_n)))$, for every $\theta(x_1, \dots, x_n) \in L(M)$, where g_i is either f or the identity. We check this by induction on n . For $n = 1$, it is the hypothesis. Assume we know it for n and take $\theta(x_1, \dots, x_{n+1}) \in L(M)$ and some g_1, \dots, g_{n+1} as above. Let $\bar{g} = (g_1, \dots, g_n)$, then the induction hypothesis says that $\bar{g}_* \mu^{(n)} = \mu^{(n)}$. Equivalently, for any Borel function χ on $S_n(M)$, we have

$$\int_{p \in S_n(M)} \chi(p) d\mu^{(n)} = \int_{p \in S_n(M)} \chi(\bar{g}_* p) d\mu^{(n)}.$$

In particular, taking $\chi(p) = \mu_x(\theta(a_1, \dots, a_n, x))$ for some (any) $(a_1, \dots, a_n) \models p$, this equation becomes

$$\mu^{(n+1)}(\theta(x_1, \dots, x_n, x_{n+1})) = \mu^{(n+1)}(\theta(g_1(x_1), \dots, g_n(x_n), x_{n+1})).$$

As $f_* \mu = \mu$, we also have $\chi(p) = \mu_x(\theta(a_1, \dots, a_n, f(x)))$ for some (any) $(a_1, \dots, a_n) \models p$ and using this expression in the right side of the equation gives

$$\mu^{(n+1)}(\theta(x_1, \dots, x_n, x_{n+1})) = \mu^{(n+1)}(\theta(g_1(x_1), \dots, g_n(x_n), f(x_{n+1}))).$$

□

Claim 2 : For every $\phi(x) \in L(\bar{M})$ and ϵ , there are only finitely many values of $t \in [0, 1]$ for which $|\nu(\phi(x_t)) - \nu(\phi(f(x_t)))| > \epsilon$.

Proof. Assume not. Then without loss we can find a sequence of reals $t_0 < t_1 < \dots \in [0, 1]$ such that $|\nu(\phi(x_{t_i})) - \nu(\phi(f(x_{t_i})))| > \epsilon$ for every even $i < \omega$. Let I be $\{t_{2i} : i < \omega\}$ and consider the measure $\eta_{x_1, x_2, \dots}$ defined by

$$\eta(\theta(x_1, \dots, x_n)) = \nu(\theta(x_{t_1}, \dots, x_{t_n}))$$

for any $\theta \in L(\bar{M})$. Define also $\eta'_{x_1, x_2, \dots}$ by

$$\eta'(\theta(x_1, \dots, x_n)) = \nu^I(\theta(x_{t_1}, \dots, x_{t_n})).$$

Then η and η' are indiscernible. By Proposition 4.2.6 and removing finitely many points from the sequence (t_i) if necessary, we may assume that for all $i < \omega$, $|\eta(\phi(x_i)) - \eta(\phi(x_{i+1}))| \leq \epsilon/4$ and same with η' instead of η . But this is a contradiction since $\eta(\phi(x_i))$ and $\eta'(\phi(x_i))$ are equal for odd i and differ by at least ϵ for even i . \square

It follows that the average of ν is the same as the one of ν^I for each I . As $f_*(\mu_\Sigma)$ is the average of ν^I for $I = [0, 1]$, we have $f_*(\mu_\Sigma) = \mu_\Sigma$. \square

As an application, we give a short proof of a result from [33], Section 7. (It is not stated explicitly there, as the notion of generically stable measure had not been introduced yet, but is the content of the pages from Lemma 7.1 to Lemma 7.6.) If G is a definable group and M a model, by a measure μ_x being $G(M)$ -invariant, we mean that μ_x concentrates on G (i.e. $\mu_x(x \in G) = 1$) and for each $g \in G(M)$ and $\phi(x)$ a formula, $\mu_x(\phi(x)) = \mu_x(\phi(g.x))$.

Proposition 4.4.5. *Let M be a model, and G a definable group. Assume there is $\mu \in \mathcal{M}_x(M)$ a $G(M)$ -invariant measure. Then μ extends to a global generically stable $G(M)$ -invariant measure.*

Proof. First, find a global extension $\tilde{\mu}$ of μ that is $G(M)$ -invariant, and M_1 -invariant for some small $M_1 \supset M$. This can be done with Keisler's smooth measure construction, see [33], Lemma 7.6. Let $\tilde{\mu}_\Sigma$ be a symmetrization of it over M_1 . Then by the previous propositions, $\tilde{\mu}_\Sigma$ is generically stable and $G(M)$ -invariant. \square

§4.5 Smooth measures

The goal of this section is to prove that generically stable measures in o-minimal theories or in the theory of the p-adics are smooth.

We start by giving a characterization of smoothness which will be useful for proving that measures are smooth. Let $M \models T$ and let $\phi(x, y) \in L(M)$. For $b \in \bar{M}$, we define the *border* of $\phi(x, b)$ (over M) as $\partial_x^M \phi(x, b) = \{p \in S_x(M) : \text{there are } a, a' \models p \text{ such that } \phi(a, b) \wedge \neg \phi(a', b) \text{ holds}\}$. This is a closed subset of the space of types $S_x(M)$. We will often omit x and M in the notation. Note that $\partial^M \phi(x, b)$ depends only on $q = \text{tp}(b/M)$, so we will also sometimes write $\partial^M \phi(x, q)$ for $\partial^M \phi(x, b)$.

Lemma 4.5.1 (Characterization of smoothness). *Let $\mu \in \mathcal{M}_x(\bar{M})$. Then μ is smooth if and only if $\mu(\partial^M \phi(x, b)) = 0$ for all $\phi(x, y) \in L(\bar{M})$ and all $b \in \bar{M}$.*

Proof. Let $\phi(x, y)$ and b be as in the statement of the lemma. Let $O \subseteq S_x(\bar{M})$ be the set of types p such that $p \vdash \phi(x, b)$. And let $F = \partial^M \phi(x, b)$. Then for any extension ν of μ , we have $\mu(O) \leq \nu(\phi(x, b)) \leq \mu(O) + \mu(F)$. Therefore if $\mu(F) = 0$ for all such ϕ and b , then μ is smooth.

Conversely, assume that μ is smooth and let $\phi(x, y)$ and b be as above. Let $\epsilon > 0$ and take $\nu_i^1(x)$, $\nu_i^2(x)$, $\psi_i(y)$, $i = 1, \dots, n$ be as in Proposition 4.2.9. Let i be such that $\psi_i(b)$ holds. Then $\partial \phi(x, b) \subseteq \nu_i^2(x) \setminus \nu_i^1(x)$. Therefore $\mu(\partial \phi(x, b)) < \epsilon$. As this is true for all $\epsilon > 0$, $\mu(\partial \phi(x, b)) = 0$. \square

To illustrate this, assume T is o-minimal, let $M \prec N$ be models of T and let $\phi(x) \in L(N)$ be a formula, x a single variable. By o-minimality, $\phi(x)$ is a finite union of (closed or open) intervals. Let a_0, \dots, a_{n-1} denote those end points that lie in $N \setminus M$. Then $\partial^M \phi = \{tp(a_k/M) : k < n\}$. In particular, it is finite.

Lemma 4.5.2. *Let $\mu \in \mathcal{M}_x(\mathcal{C})$ be a global measure, smooth over M . Let f be an M -definable function whose range is the sort of the variable x . Then $f_*(\mu)$ is smooth.*

Proof. Let $\lambda_y = f_*(\mu)$ and let $\phi(y)$ be an M -definable set. Let $F = \partial^M \phi(y)$. Define also $\psi(x) = \phi(f(x))$ and $G = \partial^M \psi(x)$. Then $\lambda(F) = \mu(G) = 0$ as μ is smooth. By Lemma 4.5.1, $f_*(\mu)$ is smooth. \square

The following easy fact will be used implicitly in what follows.

Lemma 4.5.3. *Let μ be a global smooth measure. Assume that μ is M -invariant, then μ is smooth over M .*

Proof. Proposition 4.2.9 gives us formulas $\psi_i(y, d)$, $\nu_i^1(x, d)$ and $\nu_i^2(x, d)$ with $d \in \bar{M}$ satisfying three properties as stated. It is enough to show that we can find $d' \in M$ such that the formulas $\psi_i(y, d')$, $\nu_i^1(x, d')$, $\nu_i^2(x, d')$ satisfy the same properties. Now the condition imposed on d' by the first 2 properties is clopen. By Corollary 4.2.10, μ is definable. As it is M -invariant, it must be definable over M , therefore the condition imposed by the third point is open. So we are looking for d' in some open set of $S(M)$. As we know that this set is non-empty, it must intersect the set of realized types, and we find the required d' . \square

Lemma 4.5.4. *Let S_1, S_2 be two sorts and $f : S_1 \rightarrow S_2$ a surjective definable map. Let $\mu \in \mathcal{M}_x(\bar{M})$ be a generically stable measure on the sort S_2 . Then there is $\eta \in \mathcal{M}_y(\bar{M})$ a generically stable measure on the sort S_1 such that $f_*(\eta) = \mu$.*

Proof. Let M be a small model such that μ is M -invariant. We first show that we can find a global measure ν which is M -invariant and such that $f_*(\nu) = \mu$. To see this let Ω_1 be the boolean algebra of definable sets of the form $f^{-1}(D)$, D a \bar{M} -definable set of S_2 , and Ω_2 the algebra of sets defined by a formula of the form $\phi(x, a) \Delta \phi(x, a')$ for

$\alpha, \alpha' \in \mathfrak{C}$, $\alpha \equiv_M \alpha'$. Define a partial measure ν_0 on $\Omega = \langle \Omega_1, \Omega_2 \rangle$ by $\nu_0 = f^{-1}(\mu)$ on Ω_1 and $\mu_0 = 0$ on Ω_2 . By Lemma 4.2.1, ν_0 extends to a global measure ν which satisfies the requirement.

Now consider a symmetrization η of ν over M . Then η is generically stable. Furthermore, by Proposition 4.4.4, $f_*(\eta)$ is a symmetrization of $f_*(\nu)$ over M . And as $f_*(\nu) = \mu$ is generically stable, again by Proposition 4.4.4, $f_*(\eta) = \mu$. \square

Corollary 4.5.5. *Let S_1, S_2 be two sorts and $f : S_1 \rightarrow S_2$ a surjective definable map. Assume that all generically stable measures on S_1 are smooth, then it is also the case for S_2 .*

Proof. Let μ a generically stable measure on S_2 . Take η a generically stable measure on S_1 given by Lemma 4.5.4. By hypothesis, η is smooth. So $\mu = f_*(\eta)$ is also smooth. \square

We will use this ad-hoc criterion for smoothness.

Proposition 4.5.6. *Assume T has definable Skolem functions. Let \mathcal{S} be a set of imaginary sorts containing the main sort \bar{M} . Assume that for any model N , any formula $\phi(\bar{x}, y) \in L(N)$ (y a single variable from the main sort) is a boolean combination of formulas of the form $R(f(\bar{x}), y)$ where R is a \emptyset -definable relation and f is an N -definable function taking values in a sort from \mathcal{S} . Assume that for each $S \in \mathcal{S}$, all generically stable measures over S are smooth.*

Then any generically stable measure is smooth.

Proof. Let $M \prec N$ be two models of T . Assume that for all n , all n types over M are realized in N . Let $\mu \in \mathcal{M}(N)$ be an M -invariant generically stable measure in k variables. It is enough to show that any formula of the form $\phi(\bar{x}, c)$, $c \in \bar{M}$ a 1-tuple and $\phi(\bar{x}, y) \in L(N)$, has the same measure in any extension of μ . (Because then $N(c)$ is a model of T over which μ has a unique extension, and we can replace N by it and iterate.) By hypothesis, we may assume that $\phi(\bar{x}, y) = R(f(\bar{x}), y)$ for some R and f as above.

If ν is a global extension of μ , then $\nu(\phi(\bar{x}, c)) = f_*(\nu)(R(z, c))$. Now $f_*(\nu)$ is generically stable and N -invariant. By hypothesis, it is smooth over N . Therefore $\nu(\phi(\bar{x}, c))$ is determined. \square

Corollary 4.5.7. *If T is o-minimal, then any generically stable measure is smooth.*

Proof. We will check the hypothesis of the previous proposition for $\mathcal{S} = \{\bar{M}\}$ (the main sort). Let T be o-minimal. Every formula $\phi(\bar{x}, y)$ is a boolean combination of formulas of the form $y < f(\bar{x})$ and $y = f(\bar{x})$ where f is a definable function. So the first part of the hypothesis is satisfied. Next, consider μ_x a global generically stable measure in dimension 1. Let $\phi(x)$ be a formula, x a single variable, with parameters in some extension $\bar{N} \succ \bar{M}$. As explained after the proof of Lemma 4.5.1, $\partial^{\bar{M}}\phi$ is a finite set of non-realized types. Finiteness easily implies that all the types in $\partial^{\bar{M}}\phi$ are generically stable. As there are no non-realized generically stable types in T , $\mu(\partial^{\bar{M}}\phi) = 0$. By Lemma 4.5.1, μ is smooth. Proposition 4.5.6 therefore applies. \square

Corollary 4.5.8. *Let $T = \text{Th}(\mathbb{Q}_p)$, for some p , then any generically stable measure is smooth.*

Proof. This case is similar to the o-minimal one. Let Γ denote the value group. For $1 \geq k, n < \omega$, let $\mathfrak{B}_{k,n}$ be the set of canonical parameters of definable sets of the form $\{x : \text{val}(x - a) \equiv k \pmod{n}\}$, $a \in \bar{M}$ and \mathfrak{B} be the set of canonical parameters of balls $\{x : \text{val}(x - a) = \alpha\}$ for $a \in \bar{M}$, $\alpha \in \Gamma$.

We will check that Proposition 4.5.6 applies with $S = \bigcup_{k,n} \mathfrak{B}_{k,n} \cup \{\bar{M}, \mathfrak{B}\}$.

We leave it to the reader to check that all generically stable measures in one variable from \bar{M} or from Γ are smooth (this can be done as in the o-minimal case : check that the border of a definable set is finite).

Next, let μ_x be a generically stable measure on \mathfrak{B} . Let val denote the natural map from \mathfrak{B} to Γ . Then, $\text{val}_*(\mu)$ is generically stable and therefore is smooth. We may assume that $\text{val}_*(\mu)$ is either a realized type or an atomless measure. Assume $\text{val}_*(\mu) = "x = \alpha"$ for some $\alpha \in \Gamma$. Then μ is a measure concentrating on \mathfrak{B}_α : the sort of balls of radius α . There is a surjective map $\pi : \bar{M} \rightarrow \mathfrak{B}_\alpha$, so by Corollary 4.5.5, μ is smooth. Now if $\text{val}_*(\mu)$ is atomless, one can check by inspection that $\mu(\partial\phi(x)) = 0$ for every definable set $\phi(x)$.

Finally, for any $k, n < \omega$, there is a surjective map from \mathfrak{B} to $\mathfrak{B}_{k,n}$, so all generically stable measures there are also smooth.

Let $\phi(x) \in L(A)$ be a definable set in dimension 1. Then $\phi(x)$ can be written as a finite boolean combination of formulas of the form $x \in b$ with b in some $\mathfrak{B}_{k,n}$, $k, n < \omega$ or in \mathfrak{B} . We can choose the decomposition such that b is A -definable. Therefore Proposition 4.5.6 applies. \square

Theories in which all generically stable measures are smooth will be studied in a subsequent work [61], where equivalent characterizations will be given along with some properties. In particular, it will be shown that in a dp-minimal theory with no generically stable type in the main sort, all generically stable measures are smooth, generalizing the two corollaries above.

§4.6 σ -additive measures

Recall that if $M \models T$, an externally definable subset of M is a subset of the form $\phi(M)$ where $\phi(x) \in L(\bar{M})$. Assume that the model M is equipped with a σ -algebra \mathcal{A} such that externally definable sets are measurable. Then if μ_0 is a (σ -additive) probability measure on (M, \mathcal{A}) , μ_0 induces a global Keisler measure μ . Namely $\mu(\phi(x)) = \mu_0(\phi(M))$, for $\phi(x) \in L(\bar{M})$.

Theorem 4.6.1. *Let T be NIP, $M \models T$ equipped with a σ -algebra \mathcal{A} . Assume that any externally definable subset of M^2 is measurable for the product algebra $\mathcal{A}^{\otimes 2}$. Let λ be a probability measure on (M, \mathcal{A}) . Then λ induces a global Keisler measure μ and μ is generically stable.*

Proof. Note first that by construction μ is finitely satisfiable in M (hence also M -invariant).

We have at our disposal two different amalgams of μ_x by μ_y . The first one is $\mu_{xy}^{(2)} = \mu_x \otimes \mu_y$ from the model-theoretic world. The second one comes from probability theory : we may form the product measure $\lambda^2 = \lambda \times \lambda$ which is a σ -additive measure on $(M^2, \mathcal{A}^{\otimes 2})$. By hypothesis, λ^2 induces a global Keisler measure μ_{xy}^2 . Of course, $\mu_{xy}^{(2)}$ and μ_{xy}^2 coincide on products $\phi(x) \wedge \psi(y)$ (they are both separated amalgam of μ_x and μ_y). We will prove that in fact $\mu_{xy}^2 = \mu_{xy}^{(2)}$. For this, it is enough to check that μ_{xy}^2 is a fs extension in M of μ_x .

Let $\theta(x, y), \phi(x) \in L(\bar{M})$. We have :

$$\begin{aligned} \mu_{xy}^2(\theta(x, y) \wedge \phi(x)) &= \int_{(a, b) \in M^2} \theta(a, b) \wedge \phi(a) d\lambda^2 \\ &= \int_{a \in \phi(M)} \int_{b \in M} \theta(a, b) d\lambda d\lambda \\ &\leq \lambda(\phi(M)) \cdot \sup_{a \in \phi(M)} \int_{b \in M} \theta(a, b) d\lambda \\ &= \mu_x(\phi(x)) \cdot \sup_{a \in \phi(M)} \mu_y(\theta(a, y)). \end{aligned}$$

By Proposition 4.3.5, this proves that $\mu_{xy}^2 = \mu_x \otimes \mu_y$.

By the usual Fubini theorem, λ^2 is a symmetric measure. Therefore it is also the case for $\mu_{xy}^{(2)}$: $\mu_{xy}^{(2)}(\theta(x, y)) = \mu_{xy}^{(2)}(\theta(y, x))$ for all $\theta(x, y) \in L(\bar{M})$. By Theorem 4.2.7, this implies that μ_x is generically stable. \square

Remark 4.6.2. The assumption that externally definable sets are measurable for the product σ -algebra is of course necessary. Consider for example an ω_1 -saturated model M of RCF. Let $p \in S(M)$ be the type at $+\infty$ and \tilde{p} the global co-heir of p . Then \tilde{p} is induced by a σ -additive measure on M (equipped with the Borel σ -algebra). It is not generically stable, and note that the set $\{(x, y) \in M^2 : x \leq y\}$ is not measurable for the product algebra.

As a corollary, we recover the following result from [29], Section 6.

Corollary 4.6.3. *Let M be either \mathbb{R} : the standard real numbers equipped with any o-minimal structure expanding the field operations, or \mathbb{Q}_p : the standard p -adic field. Let μ_0 be a σ -additive measure on M , then μ_0 induces a smooth Keisler measure μ .*

Proof. Theorem 4.6.1 implies that μ is generically stable, then using Corollary 4.5.7 or 4.5.8, we deduce that it is smooth. \square

Question 4.6.4. *More generally, if we assume that \mathcal{A} is generated as a σ -algebra by definable sets, is it the case that μ is smooth ?*

Proposition 4.6.5. *Let $M = \mathbb{R}$ be an o-minimal expansion of the standard model. Let μ be a smooth measure over M^k concentrating on $[0, 1]^k$. Then there is a σ -additive Borel measure λ on $[0, 1]^k$ such that λ induces μ as a Keisler measure.*

Proof. Let Ω_k be the subspace of $S_k(M)$ of types that concentrate on $[0, 1]^k$. Note that the standard part application st induces an application $st : \Omega_k \rightarrow [0, 1]^k$. This is a Borel map (as the inverse image of a closed set is closed). In particular, we can consider the pushforward measure $\lambda = st_*(\mu)$. It is a σ -additive measure on $[0, 1]^k$. We now show that λ induces μ as a Keisler measure, *i.e.*, that $\mu(X) = \lambda(X)$ for any definable set X .

Let X be a definable set of $[0, 1]^k$. Assume first that X is closed. We have a definable map $d_X : M^k \rightarrow M$ such that $d_X(\bar{x})$ is the distance of \bar{x} to X . For $\epsilon > 0$, let X_ϵ be the closed ϵ -neighborhood of X : $X_\epsilon = \{\bar{x} \in [0, 1]^k : d_X(\bar{x}) \leq \epsilon\}$. It is also a definable set. We have $st^{-1}(X) = \bigcap_{\epsilon > 0} X_\epsilon$. Let p be a type in $st^{-1}(X) \setminus X$. Then $(d_X)_*(p)$ is the type 0^+ of $S(M)$. As μ is smooth, $(d_X)_*(\mu)$ is also smooth and it is not possible that it has an atom on 0^+ . Thus $\mu(st^{-1}(X) \setminus X) = 0$ and $\mu(X) = \lambda(X)$.

We treat the general case by induction on the dimension of X . The case of dimension 0 is trivial. So let X be any definable set. Let O be its interior and \bar{X} its closure. Then $D = \bar{X} \setminus O$ has lower dimension than X . We know that $\mu(\bar{X}) = \lambda(\bar{X})$ and by induction $\mu(\bar{X} \setminus X) = \lambda(\bar{X} \setminus X)$. Hence $\mu(X) = \lambda(X)$. \square

§5.1 Introduction

We study one way in which stability and order can interact in an NIP theory. More precisely, we are interested in the situation where stability and order are intertwined. We start by giving some very simple examples illustrating what we mean.

Consider $M_0 \models \text{DLO}$. A type of $S_1(M_0)$ is determined by a cut in M_0 and two types corresponding to different cuts are orthogonal. If we take now M_1 a model of some o-minimal theory, still a 1-type is determined by a cut, but in general, types that correspond to different cuts are not orthogonal. However this is true over indiscernible sequences in the following sense : assume $\langle a_t : t < \omega + \omega \rangle \subset M_1$ is an indiscernible sequence. By NIP, the sequences of types $\langle \text{tp}(a_t/M_1) : t < \omega \rangle$ and $\langle \text{tp}(a_{\omega+t}/M_1) : t < \omega \rangle$ converge in $S(M_1)$. Then the two limit types are orthogonal (this follows from dp-minimality, see 5.2.28). An indiscernible sequence with that property will be called *distal*¹. A theory is distal if all indiscernible sequences are distal. So any o-minimal theory is distal.

Distality for an indiscernible sequence can be considered as an opposite notion to that of total indiscernibility.

Let now M_2 be a model of ACVF (or any other C-minimal structure) and consider an indiscernible sequence $(a_i)_{i < \omega}$ of 1-tuples. Two different behaviors are possible : either the sequence is totally indiscernible, this happens if and only if $\text{val}(a_i - a_j) = \text{val}(a_{i'} - a_{j'})$ for all $i \neq j, i' \neq j'$, or the sequence is distal. Again, this will follow from the results in Section 2, but could be proved directly. So M_2 is neither stable nor distal ; the two phenomena exist but do not interact in a single indiscernible sequence of points.

Consider now a fourth structure (a ‘colored order’) M_3 in the language $L_3 = \{\leq, E\}$: M_3 is totally ordered by \leq and E defines an equivalence relation, each E class being dense

1. Thanks to Itay Kaplan for suggesting the name.

co-dense with respect to \leq . Now an indiscernible sequence of elements from different E classes is neither totally indiscernible nor distal. Given two limit types p_x and q_y of different cuts in such a sequence, the type $p_x \cup q_y$ is consistent with xEy and with $\neg xEy$. Here it is clear that the ‘stable part’ of a type should be its E-class.

The idea behind the work in this paper is that every ordered indiscernible sequence in an NIP theory should look like a colored order : there is an order for which different cuts are orthogonal and a something stable on top of it which does not see the order (see Section 3).

A word about measures

Keisler measures will be used a little in this work, however the reader not familiar with them can skip all parts referring to measures without harm. For this reason, we will be very brief in recalling some facts about them and refer the reader to [33] and [29]. They however give some understanding of the intuition behind some definitions and results. We explain this now.

A Keisler measure (or simply a measure) is a Borel probability measure on a type space $S_x(A)$. Basic definitions for types (non-forking, invariance, coheir, Morley sequence etc.) generalize naturally to measures (see [33] and [29]). Of interest to us is the notion of *generically stable measure*. A measure is generically stable if it is both definable and finitely satisfiable over some small set. Equivalently, its Morley sequence is totally indiscernible. Such measures are defined and studied by Hrushovski, Pillay and the author in [29]. Furthermore, it is shown in [62] that some general constructions give rise to them, and in this sense they are better behaved than the more natural notion of generically stable type.

This paper can be considered as an attempt to understand where generically stable measures come from. What stable phenomena do generically stable measures detect ? What does the existence of generically stable measures in some particular theory tell us about types ? The first test question was : Can we characterize theories which have non-trivial generically stable measures ? Here “non-trivial” means “non-smooth” : a measure is smooth if it has a unique extension to any bigger set of parameters. This question is answered in Section 2 : a theory has a non-smooth generically stable measure if and only if it is not distal.

The main tool at our disposal to link measures to indiscernible sequences is the construction of an average measure of an indiscernible segment (see [29] Lemma 3.4 or [62] Section 3 for a more elaborate construction). Such a measure is always generically stable. The intuition we suggest is that the ‘order’ component of the sequence is evened out in the average measure and only the ‘stable’ component remains.

Organization of the paper and main results

The paper is organized as follows. The first section contains some basic facts about NIP theories and Keisler measures. We give a number of definitions concerning indiscernible sequences and some basic results illustrating how we can manipulate them.

Section 2 studies distal theories. They are defined as theories in which every indiscernible sequence is distal, as explained above. We show that this condition can also be seen through invariant types and generically stable measures. The main results can be summarized by the following theorem.

Theorem 5.1.1. *The following are equivalent :*

- T is distal,
- Any two invariant types that commute are orthogonal,
- All generically stable measures are smooth.

Furthermore, it is enough to check any one of those conditions in dimension 1.

As a consequence, o-minimal theories and the p-adics are distal as are more generally any dp-minimal theory with no generically stable type.

Section 3 can be read almost independently of the previous one : it contains a study of the intermediate case of an NIP theory that is neither stable nor distal. We deal with the problem of understanding the ‘stable part’ (or the ‘non-distal’) part of a type. We define a notion of s -domination (s for *stable*), first for points inside an indiscernible sequence and then for any point over some $|T|^+$ -saturated base M . The intuition is that if a_* s -dominates a over M , then the stable part of a_* dominates the stable part of a over M . We then define a notion of s -independence denoted $a \perp_M^s b$ which says (intuitively) that the stable parts of a and b are independent. This is a symmetric notion and is implied by forking-independence. Also, it has bounded weight. We use it to show that two commuting types behave with respect to each other like types in a stable theory (we recover some definability and uniqueness of the non-forking extension). Note that in a distal theory, all those notions are trivial.

As an application, we prove the following ‘finite-co-finite theorem’ (Theorem 5.3.30) and give an application of it to the study of externally definable sets.

Theorem 5.1.2 (Finite-co-finite theorem). *Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite. Assume that $I_1 + I_3$ is A -indiscernible and take $\phi(x; a) \in L(A)$, then the set $B = \{b \in I_2 : \models \phi(b; a)\}$ is finite or co-finite.*

The last section defines a class of theories (called *sharp*) in which the stable part of types is witnessed by generically stable types. More precisely, over a $|T|^+$ -saturated model M , every tuple is s -dominated by the realization of a generically stable type. We give a criterion for sharpness which only involves looking at indiscernible sequences of 1-tuples. In particular, any dp-minimal theory is sharp.

Our Bible concerning NIP theories are Shelah’s papers [55], [56], [53], [58] and [57]. We will however use ideas only from the first two. All the basic insights about indiscernible sequences were taken from there (although the important result on shrinking indiscernible sequences originates in [5]).

In fact, we realized after having done most of this work that the idea of ‘domination’ for indiscernible sequences was already in Shelah’s work : in Section 2 of [56] in a slightly different wording and with a very different purpose. The main additional ingredient in

Section 3 is the external characterization of domination (5.3.7) which allows us to say something about points outside of the indiscernible sequence and then to generalize to the invariant type setting.

An important property of stable theories sometimes referred to as the *Shelah reflection principle* says roughly that non-trivial relationships between a realization of a type p and some other point are reflected inside realizations of p . Internal concepts (only considering realizations of p) often imply external properties (involving the whole structure). For example regularity implies weight one. There is some evidence now that this principle is already true in NIP. See [16] for an example (weak stable embeddedness).

In this paper we will use this principle on indiscernible sequences : a property involving only the indiscernible sequence itself or extensions of it usually implies properties of the indiscernible sequence with respect to points outside (the same way total indiscernibility implies that the trace of every definable set is finite or co-finite). See Lemma 5.2.7 and Proposition 5.3.7.

Acknowledgment

I would like to thank Elisabeth Bouscaren for having patiently listened to the expositions of many erroneous proofs of results here and for helping finding mistakes. Also thanks to Udi Hrushovski for interesting remarks on earlier versions of this work and for his encouragements to pursue it.

5.1.1 Preliminaries

Basic things

We list here some terminology and basic properties of NIP theories that we will need.

We will often denote sequences of tuples by I, J, \dots . Index sets of families or sequence might be named $\mathcal{I}, \mathcal{J}, \dots$.

Recall that a theory is NIP if every indiscernible sequence I has a limit type $\lim(I/A)$ over any set A of parameters.

Assumption : From now on, until the end of the paper, we work in a NIP theory T , in a language L .

Definition 5.1.3. Two types p_x, q_y over the same domain A are *orthogonal* if $p_x \cup q_y$ defines a complete type in two variables over A .

If M is a κ -saturated model, $A \subseteq M$ satisfying $|A| < \kappa$, a type $p \in S(M)$ is A -invariant if for $a \models p$ and any tuples $b, b' \in M$, $b \equiv_A b' \rightarrow ba \equiv b'a$. We will sometimes say simply that p is an invariant type, without specifying A . Note that an invariant type has a natural extension to any larger set $B \supset M$ that we will denote by $p|_B$.

Let \mathcal{J} be a linear order. A Morley sequence indexed by \mathcal{J} of an invariant type p over some $B \supseteq A$ is a sequence $(a_t)_{t \in \mathcal{J}}$ such that $a_t \models p|_{B \cup a_{<t}}$ for every t . All Morley

sequences of \mathbf{p} over B indexed by \mathcal{J} are B -indiscernible and have the same type over B ; when $B = M$, we will denote that type by $\mathbf{p}^{(\mathcal{J})}$.

If \mathbf{p}_x and \mathbf{q}_y are two types over M and \mathbf{p} is invariant, we can define the product $\mathbf{p}_x \otimes \mathbf{q}_y$ as the element of $S_{xy}(M)$ defined as $\text{tp}(\mathbf{a}, \mathbf{b}/M)$ where $\mathbf{b} \models \mathbf{q}_y$ and $\mathbf{a} \models \mathbf{p}_x|_{M\mathbf{b}}$. If \mathbf{q} is also an invariant type, then $\mathbf{p}_x \otimes \mathbf{q}_y$ is invariant. In this case, we can also build the product $\mathbf{q}_y \otimes \mathbf{p}_x$. When the two products are equal, we say that \mathbf{p} and \mathbf{q} commute.

Note that \otimes is associative. In particular if \mathbf{p} and \mathbf{q} commute with \mathbf{r} , then \mathbf{r} commutes with $\mathbf{p} \otimes \mathbf{q}$.

Recall the notion of generically stable type from [55] and [33] : an invariant type $\mathbf{p} \in S(M)$ is generically stable if it is both definable and finitely satisfiable in some small model $N \subset M$. Equivalently, its Morley sequence is totally indiscernible.

Measures

As we mentioned in the introduction, we will not recall all definitions concerning measures. Instead, we refer the reader to [33] and [29]. The latter paper contains in particular the definition of a generically stable measure. Also the introduction of [62] contains a concise account of the definitions and basic results we will need, but without proofs.

We however recall the following from [62] :

A measure $\mu \in \mathcal{M}_x(M)$ is *smooth* if it has a unique extension to any $N \supset M$. For any formula $\phi(x, d)$, $d \in \mathcal{C}$, let $\partial_M \phi$ denote the closed set of $S_x(M)$ of types \mathbf{p} such that there are \mathbf{a}, \mathbf{a}' two realizations of \mathbf{p} satisfying $\phi(\mathbf{a}, d) \wedge \neg \phi(\mathbf{a}', d)$.

Fact 5.1.4 (Lemma 4.1 of [62]). *The measure $\mu \in \mathcal{M}_x(M)$ is smooth if and only if $\mu(\partial_M \phi) = 0$ for all formulas $\phi(x, d)$, $d \in \mathcal{C}$.*

Indiscernible sequences and cuts

The notation $I = I_1 + I_2$ means that the sequence I is the concatenation of the sequences I_1 and I_2 : I_1 is an initial segment of I and I_2 the complementary final segment. This operation is associative, and we will also use it to denote the concatenation of three or more sequences. It may be the case that one of the sequences is finite. In particular, when \mathbf{b} is a tuple, we may write $I_1 + \mathbf{b} + I_2$ to denote $I_1 + \langle \mathbf{b} \rangle + I_2$ where $\langle \mathbf{b} \rangle$ is the sequence of length 1 whose only member is \mathbf{b} .

If $I = I_1 + I_2$, we will say that (I_1, I_2) is a *cut* of I .

By the EM-type (over A) of an indiscernible sequence $I = \langle \mathbf{a}_i : i \in \mathcal{I} \rangle$, we mean the family $(\mathbf{p}_n)_{n < \omega}$, where $\mathbf{p}_n \in S_n(A)$ is the type of $(\mathbf{a}_{\sigma(k)})_{k < n}$ for $\sigma : n \rightarrow \mathcal{I}$ any increasing embedding.

We now introduce a number of definitions that will be useful for handling indiscernible sequences.

Definition 5.1.5 (Cuts). If $J \subset I$ is a convex subsequence, a cut $\mathbf{c} = (I_1, I_2)$ is said to be *interior* to J if $I_1 \cap J$ and $I_2 \cap J$ are infinite.

A cut is *Dedekind* if both I_1 and I_2^* (I_2 with the order reversed) have infinite cofinality.

If $\mathbf{c} = (I_1, I_2)$ and $\mathbf{d} = (J_1, J_2)$ are two cuts of the same sequence I , then we write $\mathbf{c} \leq \mathbf{d}$ if $I_1 \subseteq J_1$.

We write $(I'_1, I'_2) \leq (I_1, I_2)$ if I'_1 is an end segment of I_1 and I'_2 an initial segment of I_2 . A *polarized cut* is a pair $(\mathbf{c}, \varepsilon)$ where \mathbf{c} is a cut (I_1, I_2) and $\varepsilon \in \{1, 2\}$ is such that I_ε is infinite. We will write the polarized cut \mathbf{c}^- if $\varepsilon = 1$ and \mathbf{c}^+ if $\varepsilon = 2$.

Given a polarized cut $\mathbf{c}^\bullet = ((I_1, I_2), \varepsilon)$ and a set A of parameters, we can define the *limit type* of \mathbf{c}^\bullet denoted by $\lim(\mathbf{c}^\bullet/A)$ as the limit type of the sequence I_1 or I_2^* depending on the value of ε .

If a cut \mathbf{c} has a unique polarization, or if we know both polarizations give the same limit type over A , we will write simply $\lim(\mathbf{c}/A)$.

If $\mathbf{c} = (I_1, I_2)$ is a cut, we say that the tuple \mathbf{b} *fills* the cut \mathbf{c} if $I_1 + \mathbf{b} + I_2$ is indiscernible. Similarly, if $\bar{\mathbf{b}}$ is a sequence of tuples, we will say that $\bar{\mathbf{b}}$ fills \mathbf{c} if the concatenation $I_1 + \bar{\mathbf{b}} + I_2$ is indiscernible.

The following definition is from [55].

Definition 5.1.6. Let $\mathbf{c} = (I_1, I_2)$ be a Dedekind cut. A set A *weakly respects* \mathbf{c} if $\lim(\mathbf{c}^+/A) = \lim(\mathbf{c}^-/A)$. It *respects* \mathbf{c} if for every finite $A_0 \subseteq A$, there is I'_1 cofinal in I_1 and I'_2 coinital in I_2 such that $I'_1 + I'_2$ is indiscernible over A_0 .

Note that $\lim(\mathbf{c}^\bullet/\mathcal{C})$ is an invariant type, in fact finitely satisfiable over the sequence I . We will simply denote it by $\lim(\mathbf{c}^\bullet)$.

If \mathbf{c}_1 and \mathbf{c}_2 are two distinct polarized cuts in an indiscernible sequence I then $\lim(\mathbf{c}_1)$ and $\lim(\mathbf{c}_2)$ commute : $\lim(\mathbf{c}_1)_x \otimes \lim(\mathbf{c}_2)_y = \lim(\mathbf{c}_2)_y \otimes \lim(\mathbf{c}_1)_x$. More precisely $\phi(x, y) \in \lim(\mathbf{c}_1)_x \otimes \lim(\mathbf{c}_2)_y$ if and only if for some J_1 cofinal in \mathbf{c}_1 and J_2 cofinal in \mathbf{c}_2 , $\phi(a, b)$ holds for $(a, b) \in J_1 \times J_2$.

Definition 5.1.7 (Polycut). A polycut is a sequence $(\mathbf{c}_i)_{i \in J}$ of pairwise distinct cuts.

The definitions given for cuts extend naturally to polycuts : a polarized polycut is a family of polarized cuts. If $\mathbf{c} = (\mathbf{c}_i)_{i \in J}$ is a polarized polycut, then we define $\lim(\mathbf{c}) = \bigotimes_{i \in J} \lim(\mathbf{c}_i)$. It is a type in variables $(x_i)_{i \in J}$. A tuple $(a_i)_{i \in J}$ *fills* \mathbf{c} if the sequence I with all the points a_i added in their respective cut is indiscernible. Note that this is stronger than asking that each a_i fills \mathbf{c}_i .

Definition 5.1.8 (I-independent). Let I be a dense indiscernible sequence, $\mathbf{c}_1, \dots, \mathbf{c}_n$ pairwise distinct cuts in I and a_1, \dots, a_n filling those cuts, then a_1, \dots, a_n are independent over I (or I -independent) if the tuple (a_1, \dots, a_n) fills the polycut $(\mathbf{c}_1, \dots, \mathbf{c}_n)$.

We will use the notation $\mathbf{a} \perp_I \mathbf{b}$ to mean that \mathbf{a} and \mathbf{b} are independent over I , *i.e.*, that $I \cup \{\mathbf{a}\} \cup \{\mathbf{b}\}$ remains indiscernible (where $I \cup \{\mathbf{a}\} \cup \{\mathbf{b}\}$ is ordered so that \mathbf{a} and \mathbf{b} fall in their respective cuts). Note that this is a symmetric notion.

The proofs in this paper will involve a lot of constructions with indiscernible sequences. We list here the basic results and ideas we will need for that. We tried to

encapsulate in lemmas some constructions that we will use often. However, in some cases, the lemmas will not fit exactly our needs. The reader should therefore bear in minds the principles of those constructions more than the statements themselves. The constructions are grouped three parts : shrinking, expanding and sliding.

5.1.2 Shrinking

We start with the very important results concerning shrinking of indiscernibles. We give the statement as in [55, Section 3]. See also [2].

Definition 5.1.9. A finite convex equivalence relation on \mathcal{I} is an equivalence relation \sim on \mathcal{I} which has finitely many classes, all of which are convex subsets of \mathcal{I} .

Proposition 5.1.10 (Shrinking indiscernibles). *Let A be any set of parameters and $(a_t)_{t \in \mathcal{I}}$ be an A -indiscernible sequence. Let d be any tuple. Let $\phi(x_d; y_0, \dots, y_{n-1}, t)$ be a formula. There is a finite convex equivalence relation \sim on \mathcal{I} such that given :*

- $t_0 < \dots < t_{n-1}$ in \mathcal{I} ;
- $s_0 < \dots < s_{n-1}$ in \mathcal{I} with $t_k \sim s_k$ for all k ;
- $b \in A^{|t|}$,

we have $\phi(d; a_{t_0}, \dots, a_{t_{n-1}}, b) \leftrightarrow \phi(d; a_{s_0}, \dots, a_{s_{n-1}}, b)$.

Furthermore, there is a coarsest such equivalence relation.

Often we will apply this with $A = \emptyset$, in which case b does not appear.

We elaborate a little bit on this statement. We fix some parameter set A , sequence I , tuple d and formula $\phi(x_d; y_0, \dots, y_{n-1}, t)$ such that I is indiscernible over A . Consider the coarsest equivalence relation \sim satisfying the conclusion of Proposition 5.1.10.

The relation \sim induces a partition of the sequence \mathcal{I} into finitely many convex sets $\mathcal{I} = \mathcal{I}_1 + \dots + \mathcal{I}_T$ defined such that two points t and s are \sim -equivalent if and only if they belong to the same \mathcal{I}_k . We define also the corresponding partition of I as $I = I_1 + \dots + I_T$.

The $T - 1$ cuts $(I_1 + \dots + I_{k-1}, I_k + \dots + I_T)$, for $k < T$, will be called the cuts *induced by (d, ϕ) on I* (over A). For the purpose of this section, we will denote them by $\text{cut}_I(d, \phi; 0) < \dots < \text{cut}_I(d, \phi; T - 1)$. Here A is implicit to simplify the notation. Let also $T_I(d, \phi) = T$ be the number of such cuts.

Let $\mathcal{F}(n, T)$ be the set of non-decreasing functions from n to T . For any $f \in \mathcal{F}(n, T)$ and $b \in A^{|t|}$, there is a truth value $\varepsilon_{d, \phi; I}(f, b)$ such that $\phi(d; a_{t_0}, \dots, a_{t_{n-1}}, b)$ has truth value $\varepsilon_{d, \phi; I}(f, b)$ for any $t_0 < \dots < t_{n-1}$ with $t_k \in I_{f(k)}$ for all $k < T$.

To summarize, the tuple d the sequence I and the set A being fixed, we have associated, to any formula $\phi(x_d; y_0, \dots, y_{n-1}, t)$ an integer $T_I(d, \phi)$, cuts $\text{cut}_i(d, \phi; I)$ for $i < T_I(d, \phi)$ and a function $\varepsilon_I(d, \phi) : \mathcal{F}(n, T_I(d, \phi)) \times A^{|t|} \rightarrow \{\top, \perp\}$. This data completely describes the type of d over IA .

We now define a notion of similarity between types. Let d and d' two tuples of the same length, and I and I' be two sequences indiscernible over the same set A of parameters. We say that $\text{tp}(d/I)$ and $\text{tp}(d'/I')$ are similar over A , if for every formula ϕ and ψ as above we have :

- the EM-types of I and I' over A are the same ;

- $T_I(\mathbf{d}, \phi) = T_{I'}(\mathbf{d}', \phi)$;
- $\epsilon_I(\mathbf{d}, \phi) = \epsilon_{I'}(\mathbf{d}', \phi)$;
- for all $i < T_I(\mathbf{d}, \phi)$, the cuts $\text{cut}_I(\mathbf{d}, \phi; i)$ and $\text{cut}_{I'}(\mathbf{d}', \phi; i)$ are either both of infinite cofinality from the left (resp. right) or both of finite cofinality from the left (resp. right);
- for all $i < T_I(\mathbf{d}, \phi)$ and $j < T_I(\mathbf{d}, \psi)$, we have $\text{cut}_I(\mathbf{d}, \phi; i) < \text{cut}_I(\mathbf{d}, \psi; j)$ if and only if $\text{cut}_{I'}(\mathbf{d}', \phi; i) < \text{cut}_{I'}(\mathbf{d}', \psi; j)$;
- there are infinitely many elements in I between the cuts $\text{cut}_I(\mathbf{d}, \phi; i)$ and $\text{cut}_I(\mathbf{d}', \psi; j)$ if and only if there are infinitely many elements in I' between the cuts $\text{cut}_{I'}(\mathbf{d}', \phi; i)$ and $\text{cut}_{I'}(\mathbf{d}', \psi; j)$; furthermore, if there are finitely many elements in both cases, then the number of those elements is the same.

Note that in particular, if $\text{tp}(\mathbf{d}/I)$ and $\text{tp}(\mathbf{d}'/I')$ are similar over A , then $\text{tp}(\mathbf{d}/A) = \text{tp}(\mathbf{d}'/A)$.

Many notions defined in this paper concerning $\text{tp}(\mathbf{d}/I)$ only depend on the similarity class of this type. So it is convenient to note ways in which we can modify I preserving this similarity class.

Lemma 5.1.11. *Let I be a dense indiscernible sequence over a set A , and \mathbf{d} a tuple, then there is $I' \subset I$ of size at most $|T| + |\mathbf{d}|$ such that $\text{tp}(\mathbf{d}/I')$ and $\text{tp}(\mathbf{d}/I)$ are similar over A .*

Proof. Let $I_0 \subset I$ be a set consisting of :

- the extremal point of the cut $\text{cut}_I(\mathbf{d}, \phi; i)$ for every (ϕ, i) for which this cut is not Dedekind;
- two points between the two cuts $\text{cut}_I(\mathbf{d}, \phi; i)$ and $\text{cut}_I(\mathbf{d}, \psi; j)$ whenever possible;
- one point above the cut $\text{cut}_I(\mathbf{d}, \phi; i)$ whenever this is possible and one point below this cut, when possible.

Thus I_0 has size at most $|T| + |\mathbf{d}|$. Take I' to be a dense subsequence of I containing I_0 and of size at most $|T| + |\mathbf{d}|$. Then $\text{tp}(\mathbf{d}/I')$ is similar to $\text{tp}(\mathbf{d}/I)$. \square

Lemma 5.1.12. *Let $I = (\mathbf{a}_t)_{t \in J}$ be A -indiscernible with J of cofinality at least $|T|^+$, then for any finite tuple \mathbf{d} , there is an end segment I' of I that is indiscernible over $A\mathbf{d}$.*

Proof. Simply take I' to be to the right of all the cuts $\text{cut}_I(\mathbf{d}, \phi; i)$. \square

5.1.3 Expanding

Let I be an indiscernible sequence over some set A , and \mathbf{d} any tuple. We now study how one can extend I to some bigger sequence I' maintaining the similarity type of $\text{tp}(\mathbf{d}/I)$ over A .

First, if I is endless, there is a limit type $\lim(I/A\mathbf{d})$ as defined above. If J realizes a Morley sequence of that type, then $I + J^*$ is indiscernible, where J^* is the sequence J with the opposite order. Also $\text{tp}(\mathbf{d}/I + J^*)$ is similar to $\text{tp}(\mathbf{d}/I)$ over A .

Consider now a cut $\mathbf{c} = (I_1, I_2)$ of I . If I_1 is endless, then we can similarly consider K a Morley sequence of $\lim(I_1)$ over IA . Then $I_1 + K^* + I_2$ is indiscernible and $\text{tp}(\mathbf{d}/I_1 + K^* + I_2)$

is similar to $\text{tp}(\mathbf{d}/I_1 + I_2)$. If I_2 has no first element, then we can similarly extend by realizing a Morley sequence in $\lim(I_2^*)$. Note that unless the cut \mathfrak{c} is induced by (\mathbf{d}, ϕ) on I for some formula ϕ , then $\lim(I_1/I\text{Ad}) = \lim(I_2^*/I\text{Ad})$.

If we want to extend the sequence I by adding elements in different cuts, we can iterate the above procedure. Note that the order in which we chose the cuts does not matter since the different limit types commute with each other.

We therefore conclude the following lemma.

Lemma 5.1.13. *Let $I = (e_i)_{i \in \mathcal{I}}$ be an indiscernible sequence over some set A . Assume \mathcal{I} is dense without endpoints. Let \mathbf{d} be any tuple and let $\mathcal{J} \supset \mathcal{I}$ be any linearly ordered set extending \mathcal{I} . Then there are tuples $(e_i)_{i \in \mathcal{J} \setminus \mathcal{I}}$ such that the sequence $J = (e_i)_{i \in \mathcal{J}}$ is indiscernible over A and $\text{tp}(\mathbf{a}/J)$ is similar to $\text{tp}(\mathbf{a}/I)$ over A .*

5.1.4 Sliding

We are now concerned with the situation where we have A , I and \mathbf{d} as above, and we want to produce some \mathbf{d}' with the same similarity type as \mathbf{d} , but such that the cuts induced by \mathbf{d}' are different from those induced by \mathbf{d} . We see this as *sliding* the point \mathbf{d} along the sequence.

We state the result in a slightly more general form involving two different sequences.

Lemma 5.1.14. *Let I, J be two dense sequences, indiscernible over some set A . Assume they have no endpoints and have the same EM-type over A . Let \mathbf{d} be any tuple. For any formula ϕ such that $\text{cut}_I(\mathbf{d}, \phi; i)$ is well defined, pick a cut $\mathfrak{d}(\phi; i)$ of J such for any ϕ , ψ , i, j for which this makes sense :*

- *the cuts $\text{cut}_I(\mathbf{d}, \phi; i)$ and $\mathfrak{d}(\phi; i)$ are either both of infinite cofinality from the left (resp. right) or both of finite cofinality from the left (resp. right) ;*
- *we have $\mathfrak{d}(\phi; i) < \mathfrak{d}(\psi; j)$ if and only if $\text{cut}_I(\mathbf{d}, \phi; i) < \text{cut}_I(\mathbf{d}, \psi; j)$;*
- *there are infinitely many elements in J between the cuts $\mathfrak{d}(\phi; i)$ and $\mathfrak{d}(\psi; j)$ if and only if there are infinitely many elements in I between the cuts $\text{cut}_I(\mathbf{d}, \phi; i)$ and $\text{cut}_I(\mathbf{d}', \psi; j)$.*

Then there is a point \mathbf{e} such that $\text{tp}(\mathbf{e}/J)$ is similar to $\text{tp}(\mathbf{d}/I)$ over A and $\text{cut}_J(\mathbf{e}, \phi; i) = \mathfrak{d}(\phi; i)$ for any ϕ and i .

Proof. This translates into finding \mathbf{e} with a prescribed type $p(\mathbf{x})$ over AJ . Let $\theta(\mathbf{x}; \bar{\mathbf{m}}) \in p(\mathbf{x})$, $\bar{\mathbf{m}} \subset J$. Also we may assume that $\theta(\mathbf{x}; \bar{\mathbf{m}})$ is a conjunction of the form

$$\bigwedge_j \phi_j^{\epsilon_j}(\mathbf{x}; \bar{\mathbf{m}}, \mathbf{b}); \quad \mathbf{b} \in A, \bar{\mathbf{m}} \in J,$$

where ϵ_j is either 0 or 1 depending on the position of the points in $\bar{\mathbf{m}}$ with respect to the cuts $\mathfrak{d}(\phi; i)$. We can find an injection $\sigma : \bar{\mathbf{m}} \rightarrow I$ such that :

- for every $\mathbf{m}_0, \mathbf{m}_1$ in $\bar{\mathbf{m}}$, if $\mathbf{m}_0 <_J \mathbf{m}_1$, then $\sigma(\mathbf{m}_0) <_I \sigma(\mathbf{m}_1)$;
- for every index j and $\mathbf{m}_0 \in \bar{\mathbf{m}}$, the relative position of $\sigma(\mathbf{m}_0)$ and the cut $\text{cut}_I(\mathbf{a}, \phi; i)$ on I is the same as that of \mathbf{m}_0 and $\mathfrak{d}(\phi; i)$.

Then σ is a partial isomorphism and $\mathbf{a} \models \bigwedge_j \phi_j(\mathbf{x}; \sigma(\bar{\mathbf{m}}))$. Therefore $\theta(\mathbf{x}; \bar{\mathbf{m}})$ is consistent and by compactness, $p(\mathbf{x})$ is consistent. \square

Corollary 5.1.15. *Let I, J be two dense sequences with no endpoints indiscernible over some set A of same EM-type over A . Let \mathbf{a} and \mathbf{b} be tuples of the same length such that $\text{tp}(\mathbf{a}/I)$ and $\text{tp}(\mathbf{b}/J)$ are similar over A . Let \mathbf{a}' be any tuple. Then there is an indiscernible sequence $J' \supseteq J$ and a tuple \mathbf{b}' such that $\text{tp}(\mathbf{b}\mathbf{b}'/J')$ is similar to $\text{tp}(\mathbf{a}\mathbf{a}'/I)$ over A .*

Proof. By expanding, we can find a sequence J' extending J such that $\text{tp}(\mathbf{b}/J')$ is similar to $\text{tp}(\mathbf{b}/J)$ and the sequence J' is indexed by a $|T|^+$ -saturated dense linear order. It is then easy to find cuts $\mathfrak{d}(\phi; i)$ in J' as in the previous lemma corresponding to the cuts $\text{cut}_I(\mathbf{a}\mathbf{a}', \phi; i)$ in a way compatible with the cuts $\text{cut}_{J'}(\mathbf{b}, \phi; i)$ over J' . Lemma 5.1.14 gives us a tuple $\mathbf{b}_0\mathbf{b}'_0$ of same length as $\mathbf{a}\mathbf{a}'$ such that $\text{tp}(\mathbf{b}_0\mathbf{b}'_0/J')$ is similar to $\text{tp}(\mathbf{a}\mathbf{a}'/I)$. By assumption on the cuts $\mathfrak{d}(\phi; i)$, we have $\text{tp}(\mathbf{b}_0/J') = \text{tp}(\mathbf{b}/J')$ so by composing by an automorphism over J' , we obtain some \mathbf{b}' as required. \square

Corollary 5.1.16. *Let I, J be two dense sequences with no endpoints indiscernible over A and of same EM-type over A . Let \mathbf{a} and \mathbf{b} be tuples of the same length such that $\text{tp}(\mathbf{a}/I)$ and $\text{tp}(\mathbf{b}/J)$ are similar over A . Let $I' \supseteq I$ be indiscernible and let \mathbf{a}' be any tuple. Then there is an indiscernible sequence $J' \supseteq J$ and a tuple \mathbf{b}' such that $\text{tp}(\mathbf{b}\mathbf{b}'/J')$ is similar to $\text{tp}(\mathbf{a}\mathbf{a}'/I')$ over A .*

Proof. Simply apply the previous corollary with \mathbf{a}' there equal to $\mathbf{a}' \cup (I' \setminus I)$ here. \square

5.1.5 Weight and dp-minimality

Let $(I_i)_{i < \alpha}$ be a family of indiscernible sequences and A a set of parameters. We say that the sequences $(I_i)_{i < \alpha}$ are *mutually indiscernible* over A if for every $i < \alpha$, the sequence I_i is indiscernible over $A \cup \{I_j : j < \alpha, j \neq i\}$.

The following observations are from [55].

Proposition 5.1.17. *Let $(I_i)_{i < |T|^+}$ be mutually indiscernible sequences (over some set A) and let \mathbf{d} be a tuple of size at most $|T|$. Then there is some $i < |T|^+$ such that I_i is indiscernible over $A\mathbf{d}$.*

Proof. Assume not, then for every $i < |T|^+$, we can find two tuples $\bar{\mathbf{a}}_i$ and $\bar{\mathbf{b}}_i$ of increasing elements from I_i and a formula $\phi_i(\mathbf{x}, \bar{\mathbf{y}})$ such that $\mathbf{d} \models \phi_i(\mathbf{x}, \bar{\mathbf{a}}_i) \wedge \neg \phi_i(\mathbf{x}, \bar{\mathbf{b}}_i)$. Removing some sequences from the family, we may assume that $\phi_i = \phi$ does not depend on i . By mutual indiscernibility, we have $\text{tp}(\mathbf{a}_i/\{I_j : j \neq i\}) = \text{tp}(\mathbf{b}_i/\{I_j : j \neq i\})$ for all $i < |T|^+$. It follows that for every $A \subseteq |T|^+$, we can find a tuple \mathbf{d}_A such that for all $i < |T|^+$, $\mathbf{d}_A \models \phi(\mathbf{x}, \bar{\mathbf{a}}_i)$ if and only if $i \in A$. This contradicts NIP. \square

Corollary 5.1.18. *Let M be some κ -saturated model, and let $(\mathbf{p}_i)_{i < |T|^+}$ be a family of pairwise commuting invariant types over M . Let $\mathbf{p} = \bigotimes_{i < |T|^+} \mathbf{p}_i$ and $(\mathbf{a}_i)_{i < |T|^+} \models \mathbf{p}$. Let also $\mathbf{q} \in S(M)$ be any type and $\mathbf{d} \models \mathbf{q}$. Then there is $i < |T|^+$ such that $(\mathbf{a}_i, \mathbf{d}) \models \mathbf{p}_i \otimes \mathbf{q}$.*

Proof. Build a Morley sequence $\langle (\mathbf{a}_i^k)_{i < |T|^+} : 0 < k < \omega \rangle$ of \mathbf{p} over everything and set $\mathbf{a}_i^0 = \mathbf{a}_i$ for each i . Commutativity implies that the sequences $(\mathbf{a}_i^k)_{k < \omega}$, $i < |T|^+$ are mutually indiscernible. The result then follows by Proposition 5.1.17. \square

Observe in particular that if q is an invariant type, taking $b \models q|[\{a_i : i < |T|^+\}]$, we obtain that there is $i < |T|^+$ such that p_i and q commute.

We will occasionally mention dp-minimal theories. They are theories for which the notion of weight suggested by Proposition 5.1.17 is equal to 1 on 1-types. This notion was introduced by Shelah in [53].

Definition 5.1.19 (Dp-minimal). An theory T is dp-minimal if it is NIP and if for every indiscernible sequence I and 1-tuple d , there is a subdivision $I = I_1 + I_2 + I_3$ into convex sets, where I_2 is either reduced to a point or empty and I_1 and I_3 are both indiscernible over d .

Equivalently, for every two mutually indiscernible sequences I and J and 1-tuple d , one of I or J is indiscernible over d .

See [63] for the proof of the equivalence and [18] for additional information.

Examples of dp-minimal theories include o-minimal and C-minimal theories and the p-adics.

§5.2 Distal theories

5.2.1 Indiscernible sequences

We now state the main definition of this paper.

Definition 5.2.1 (Distal). An indiscernible sequence I is distal if for every dense sequence J of same EM-type as I , every distinct Dedekind cuts c_1 and c_2 of J , if a fills c_1 and b fills c_2 , then $a \downarrow_J b$.

An NIP theory T is *distal* if all indiscernible sequences are distal.

Remark 5.2.2. Equivalently the two types $\lim(c_1/J)$ and $\lim(c_2/J)$ are orthogonal.

Lemma 5.2.3. *If I is dense and has two distinct Dedekind cuts c_1 and c_2 , then it is distal if and only if $\lim(c_1/I)$ and $\lim(c_2/I)$ are orthogonal (i.e., there is no need for J in the definition).*

Proof. Left to right is obvious. We show the converse. If I is not distal, then there is some dense sequence J of same EM-type, two distinct Dedekind cuts d_1 and d_2 of J , some a_1 filling d_1 and a_2 filling d_2 such that $a_1 \not\downarrow_J a_2$. Let $\phi(a_1, a_2, \bar{m})$ be a formula witnessing that, with $\bar{m} \in I$. Take a countable $J' \subseteq J$ containing \bar{m} such that a_1 and a_2 fill Dedekind cuts of J' . Replacing J by J' , we may assume that J is countable.

Then by expanding, we can find some $J_0 \supseteq J$ and an automorphism σ mapping J_0 onto I and such that the cut d_1 (resp. d_2) is mapped to c_1 (resp. c_2) and the types $\text{tp}(a_1, a_2/J)$ and $\text{tp}(a_1, a_2, J_0)$ are similar. Then, the points $\sigma(a_1)$ and $\sigma(a_2)$ fill respectively the cuts c_1 and c_2 and $\phi(\sigma(a_1), \sigma(a_2), \sigma(\bar{m}))$ holds. Therefore $\sigma(a_1) \not\downarrow_I \sigma(a_2)$ and it follows that the two limit types $\lim(c_1/I)$ and $\lim(c_2/I)$ are not orthogonal. \square

Actually, it will follow from Lemma 5.2.7 that the hypothesis that I is dense can be removed.

EXAMPLE 5.2.4. Assume I is an indiscernible sequence, f a definable function such that $f(I)$ is totally indiscernible (non constant), then I is not distal. To see this, take \mathbf{a} and \mathbf{b} in the definition such that $f(\mathbf{a}) = f(\mathbf{b})$. See 5.2.15 for a more general result.

EXAMPLE 5.2.5. In DLO, any two 1-types concentrating on different cuts are orthogonal. It is easy then to check that it is a distal theory. We will see (Corollary 5.2.28) that in fact any o-minimal theory is distal.

Lemma 5.2.6. Assume I is a dense indiscernible distal sequence, and $\mathbf{c}_0, \dots, \mathbf{c}_{n-1}$ are pairwise distinct Dedekind cuts. If for each $i < n$, \mathbf{a}_i fills \mathbf{c}_i then the family $(\mathbf{a}_i)_{i < n}$ is I -independent.

Proof. We prove it by induction on n . for $n = 2$, it is Lemma 5.2.3. Assume it holds for n and consider a family $(\mathbf{c}_i)_{i < n+1}$ and $(\mathbf{a}_i)_{i < n+1}$ as in the hypothesis. Let $I' = I \cup \{\mathbf{a}_0\}$ (where \mathbf{a}_0 is inserted in the cut \mathbf{c}_0). Each cut \mathbf{c}_i naturally induces a cut \mathbf{c}'_i of I' . By the case $n = 2$, for each $0 < i < n + 1$, \mathbf{a}_i fills \mathbf{c}'_i . The sequence I' is also distal, so by induction $(\mathbf{a}_i)_{0 < i < n+1}$ is I' -independent. Therefore $(\mathbf{a}_i)_{i < n+1}$ is I -independent. \square

Lemma 5.2.7 (External characterization of distality). *A sequence I is distal if and only if the following property holds : For every set A , tuple \mathbf{b} and A -indiscernible sequence $I' = I_1 + I_2$ (I_1 and I_2 without endpoints, $EM\text{-}tp(I') = EM\text{-}tp(I)$), if $I_1 + \mathbf{b} + I_2$ is indiscernible, it is A -indiscernible.*

Proof. Assume that I is distal, but the conclusion does not hold. Then there is some $I' = I_1 + I_2$ and formula $\phi(\mathbf{x})$ with parameters from $A \cup I_1 \cup I_2$ which witnesses it. This means $\phi(\mathbf{b})$ holds and there is $(I'_1, I'_2) \trianglelefteq (I_1, I_2)$ such that $\neg\phi(\mathbf{a})$ holds for $\mathbf{a} \in I'_1 \cup I'_2$. Restricting even more if necessary, we may assume that $I'_1 + I'_2$ is indiscernible over the parameters of ϕ . So replacing I' by that latter sequence, we may assume that all the parameters are from A . Then, we may freely enlarge I' , so assume that it is dense.

As I' is A -indiscernible, for every cut \mathbf{c} of I' , there is \mathbf{b}' filling it such that $\phi(\mathbf{b}')$ holds. Fix an increasing sequence $(\mathbf{c}_k)_{k < \omega}$ of such cuts. For every $k < \omega$, let \mathbf{b}_k fill \mathbf{c}_k such that $\phi(\mathbf{b}_k)$ holds. The sequence I' is distal (because I' and I have same EM-type) so by Lemma 5.2.6, the sequence formed by adding all those points to I' is still indiscernible. Therefore $\phi(\mathbf{x})$ has infinite alternation number, contradicting NIP.

The converse is easy. \square

The following technical lemma will be used repeatedly.

Lemma 5.2.8 (Strong base change). *Let I be an indiscernible sequence and $A \supseteq I$ a set of parameters. Let $(\mathbf{c}_i)_{i < \alpha}$ be a sequence of pairwise distinct polarized Dedekind cuts in I . For each $i < \alpha$ let \mathbf{d}_i fill the cut \mathbf{c}_i . Then there exist $(\mathbf{d}'_i)_{i < \alpha}$ such that $tp((\mathbf{d}'_i)_{i < \alpha}/I) = tp((\mathbf{d}_i)_{i < \alpha}/I)$ and for each $i < \alpha$, $tp(\mathbf{d}'_i/A) = \lim(\mathbf{c}_i/A)$.*

Proof. Assume the result does not hold. Then by compactness, we may assume that $\alpha = n$ is finite and that there is a formula $\phi(x_0, \dots, x_{n-1}) \in \text{tp}((d_i)_{i < n}/I)$ and formulas $\psi_i(x_i) \in \text{lim}(c_i/\mathfrak{m})$ for some finite $\mathfrak{m} \in A^k$ such that $\phi(x_0, \dots, x_{n-1}) \wedge \bigwedge_i \psi_i(x_i)$ is inconsistent. Let I_0 denote the parameters of ϕ , and assume $I_0 \subseteq \mathfrak{m}$.

Assume for simplicity that $n = 2$ (the proof for $n > 2$ is the same) and without loss each c_i is polarized as c_i^- . For $i = 0, 1$ take $(J_i, J'_i) \leq c_i$ such that ψ_i holds on all elements of J_i and $J_i \cup J'_i$ contains no element of I_0 . Then $J_0 + J'_0$ and $J_1 + J'_1$ are mutually indiscernible over I_0 . So for every two cuts \mathfrak{d}_0 and \mathfrak{d}_1 respectively from $J_0 + J'_0$ and $J_1 + J'_1$, we can find points e_0 and e_1 filling those cuts (even seen as cuts of I) such that $\phi(e_0, e_1)$ holds.

Take two cuts \mathfrak{d}_0 and \mathfrak{d}_1 of I such that they are respectively interior to J_0 and J_1 . Fill \mathfrak{d}_0 by e_0 and \mathfrak{d}_1 by e_1 such that $\phi(e_0, e_1)$ holds. By hypothesis, either $\neg\psi_0(e_0)$ or $\neg\psi_1(e_1)$ holds. Assume $\neg\psi_1(e_1)$ holds. Now forget about e_0 and set $I' = I \cup \{e_1\}$. Then I' is indiscernible and we take it as our new I . Set $J'_0 = J_0$ and let J'_1 be an initial segment of J_1 not containing \mathfrak{d}_1 and make the same construction. We obtain new points (e_0^1, e_1^1) that fill the cuts $\mathfrak{d}_0^1, \mathfrak{d}_1^1$ of J'_0 and J'_1 such that $\neg\psi_0(e_0^1) \vee \neg\psi_1(e_1^1)$ holds. Without loss (as we will iterate infinitely many times) again $\neg\psi_1(e_1^1)$ holds.

Iterate this ω time to obtain a sequence of points e_1^k and cuts \mathfrak{d}_1^k in J_1 such that I with all the points e_1^k added in the cuts \mathfrak{d}_1^k is indiscernible and $\neg\psi_1(e_1^k)$ holds for all n . But $\psi_1(x)$ holds for all $x \in J_1$ so ψ_1 has infinite alternation rank, contradicting NIP. \square

Corollary 5.2.9 (Base change). *The notion of being distal is stable both ways under base change : If I is A -indiscernible, then I is distal in $T(A)$ if and only if it is distal in T .*

Proof. Assume I is distal in T . Notice that the property stated in Lemma 5.2.7 is preserved under naming parameters (because we can incorporate them in the set A). This implies that I is distal in $T(A)$.

Conversely, assume I is not distal in T . Increase I to some large A -indiscernible sequence $J_1 + J_2 + J_3$ and take $\mathfrak{a}, \mathfrak{b}$ such that $J_1 + \mathfrak{a} + J_2 + J_3$ and $J_1 + J_2 + \mathfrak{b} + J_3$ are indiscernible, but $J_1 + \mathfrak{a} + J_2 + \mathfrak{b} + J_3$ is not. By strong base change, we may assume that \mathfrak{a} and \mathfrak{b} realize the limit types over A of the cuts they define. Then $J_1 + \mathfrak{a} + J_2 + J_3$ and $J_1 + J_2 + \mathfrak{b} + J_3$ are A -indiscernible, giving a counter-example to distality in $T(A)$. \square

Lemma 5.2.10. *If T is dp -minimal and I is an indiscernible sequence of 1-tuples, not totally-indiscernible, then I is distal.*

Proof. Write $I = (d_i)_{i \in \mathcal{J}}$ and assume that it is not totally indiscernible. Working over some base A if necessary, we may assume that there is a formula $\phi(x, y) \in L(A)$ which orders the sequence I and such that I is indiscernible over A . So we have $\phi(d_i, d_j) \iff i < j$.

Without loss \mathcal{J} is a dense order and can be written as $J_1 + J_2 + J_3$, the three pieces being infinite without end points. Write $I = I_1 + I_2 + I_3$ in the obvious way. Let \mathfrak{a} fill the cut $c_a = (I_1, I_2 + I_3)$ and \mathfrak{b} fill $c_b = (I_1 + I_2, I_3)$. Assume that \mathfrak{a} and \mathfrak{b} contradict distality of I . So there is a formula $\psi(x, y) \in L(AI)$ such that $\psi(\mathfrak{a}, \mathfrak{b})$ holds and witnesses $\mathfrak{a} \not\perp_I \mathfrak{b}$.

Let $\bar{d} = (d_{i_1}, \dots, d_{i_n})$ be the parameters of ψ coming from I with $i_1 < \dots < i_n$. Let s be such that exactly i_1, \dots, i_s are from \mathcal{J}_1 and t such that exactly i_{s+1}, \dots, i_t are from \mathcal{J}_2 . Let \mathcal{J}'_1 be an end segment of \mathcal{J}_1 above i_s and \mathcal{J}'_3 an initial segment of \mathcal{J}_3 below i_{t+1} .

Let $\bar{d}_1 = (d_{i_1}, \dots, d_{i_s})$ and $\bar{d}_3 = (d_{i_{t+1}}, \dots, d_{i_n})$. Consider the sequence $J = \langle d_i \wedge \bar{d}_1 \wedge \bar{d}_3 : i \in \mathcal{J}'_1 + \mathcal{J}_2 \rangle + \langle b \wedge \bar{d}_1 \wedge \bar{d}_3 \rangle + \langle d_i \wedge \bar{d}_1 \wedge \bar{d}_3 : i \in \mathcal{J}'_3 \rangle$. It is an indiscernible sequence. By dp-minimality applied to J and α , we know that J breaks into $J_1 + J_2 + J_3$, J_2 having at most one element, and such that J_1 and J_3 are indiscernible over α . Considering the formula $\phi(x, \alpha)$, we know that J_1 must be equal to $\langle d_i \wedge \bar{d}_1 \wedge \bar{d}_3 : i \in \mathcal{J}'_1 \rangle$. And then J_2 is empty and J_3 is the rest of the sequence. In particular the tuple $b \wedge \bar{d}_1 \wedge \bar{d}_3$ lies inside J_3 as do all the parameters of $\psi(x, y)$. As $\psi(\alpha, b)$ holds but $\neg\psi(\alpha, d_i)$ holds for $i \in \mathcal{J}_3$, we get a contradiction to the indiscernability of J_3 over α . \square

Lemma 5.2.11. *Let T be distal, I and J are two mutually indiscernible sequences. Let \mathfrak{c} (resp. \mathfrak{d}) be a cut in the interior of I (resp. J). Then $\lim(\mathfrak{c}/IJ)$ and $\lim(\mathfrak{d}/IJ)$ are orthogonal.*

Proof. Write $I = (a_i)_{i \in \mathcal{J}}$ and $J = (b_j)_{j \in \mathcal{J}}$. Assume the conclusion does not hold. Then there are $\alpha \models \lim(\mathfrak{c}/IJ)$ and $\beta \models \lim(\mathfrak{d}/IJ)$ and a formula $\phi(x, y) \in L(IJ)$ such that $\phi(\alpha, \beta)$ holds, but $\lim(\mathfrak{c}) \otimes \lim(\mathfrak{d}) \vdash \neg\phi(x, y)$. Let \mathcal{K} be a countable dense linear order without end points. Pick embedding $\tau_1 : \mathcal{K} \rightarrow \mathcal{J}$ and $\tau_2 : \mathcal{K} \rightarrow \mathcal{J}$ such that :

- \mathfrak{c} induces a Dedekind cut on $\tau_1(\mathcal{K})$ and \mathfrak{d} induces a Dedekind cut on $\tau_2(\mathcal{K})$;
- identifying $\tau_1(\mathcal{K})$ and $\tau_2(\mathcal{K})$, those two Dedekind cuts are distinct;
- the parameters of $\phi(x, y)$ belong to $\{a_i : i \in \tau_1(\mathcal{K})\} \cup \{b_j : j \in \tau_2(\mathcal{K})\}$.

Let K be the sequence $\langle a_{\tau_1(t)} \wedge b_{\tau_2(t)} : t \in \mathcal{K} \rangle$. Let \mathfrak{c}' and \mathfrak{d}' denote the two cuts naturally induced by \mathfrak{c} and \mathfrak{d} on K . There are tuples b_* and a_* such that $a_* \wedge b_*$ fills \mathfrak{c}' and $a_* \wedge b$ fill \mathfrak{d}' . By distality of K , $a_* \wedge b_* \perp_K a_* \wedge b$ and $\phi(a, b)$ holds. This contradicts the assumption. \square

Definition 5.2.12 (Weakly linked). Let $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$ be an indiscernible sequence of pairs. We say that $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are weakly linked if for every disjoint subsets \mathcal{J}_1 and \mathcal{J}_2 of \mathcal{J} , $(a_i)_{i \in \mathcal{J}_1}$ and $(b_i)_{i \in \mathcal{J}_2}$ are mutually indiscernible.

Observation 5.2.13. 1. *If $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$ is A -indiscernible and $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are mutually indiscernible, then they are mutually indiscernible over A .*

2. *If $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$ is A -indiscernible and $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are weakly linked, then they are weakly linked over A .*

Lemma 5.2.14. *Let $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$ be indiscernible.*

1. *If $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are weakly linked and $(a_i)_{i \in \mathcal{J}}$ is distal, then $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are mutually indiscernible.*
2. *If $(b_i)_{i \in \mathcal{J}}$ is totally indiscernible, then $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are weakly linked.*

Proof. (1). Without loss, we may assume that \mathcal{J} is dense. Pick some finite $\mathcal{J}_2 \subset \mathcal{J}$. Then $(a_i)_{i \notin \mathcal{J}_2}$ is indiscernible over $B = (b_i)_{i \in \mathcal{J}_2}$. By applying repeatedly Lemma 5.2.7, we obtain that $(a_i)_{i \in \mathcal{J}}$ is indiscernible over B . This is enough.

(2). Assume \mathcal{J} is dense and big enough, take $\mathcal{J}_1 \subset \mathcal{J}$ finite and let $A = (a_i)_{i \in \mathcal{J}_1}$. By shrinking of indiscernibles and using total indiscernability of $(b_i)_{i \in \mathcal{J}}$, there is $\mathcal{J}_2 \subset \mathcal{J}$ of size at most $|\mathcal{T}|$ such that $(b_i)_{i \in \mathcal{J} \setminus \mathcal{J}_2}$ is indiscernible over A . By indiscernability of $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$, we may take $\mathcal{J}_2 = \mathcal{J}_1$. Therefore $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are weakly linked. \square

Corollary 5.2.15. *Let $\langle (a_i, b_i) : i \in \mathcal{J} \rangle$ be an indiscernible sequence. Assume $(a_i)_{i \in \mathcal{J}}$ is totally indiscernible and $(b_i)_{i \in \mathcal{J}}$ is distal, then $(a_i)_{i \in \mathcal{J}}$ and $(b_i)_{i \in \mathcal{J}}$ are mutually indiscernible.*

5.2.2 Invariant types

We prove here a characterization of distality in terms of invariant types.

If M is a κ -saturated model, by an invariant type over M , we mean a type $p \in S(M)$ invariant over some $A \subset M$, $|A| < \kappa$. If p and q are two invariant types over M , then we can define the products $p_x \otimes q_y$ and $q_y \otimes p_x$ as explained in the introduction. The types p and q *commute* if those two products are equal.

Lemma 5.2.16. *Assume \mathcal{T} is distal. Let M be κ -saturated and let $p, q \in S(M)$ be invariant types. If $p_x \otimes q_y = q_y \otimes p_x$, then p and q are orthogonal.*

Proof. Let $b \models q$ and let $N \prec M$ a model of size $< \kappa$ such that p and q are N -invariant. Let $I \subset M$ be a Morley sequence of p over N . Let a realize p , and build I' a Morley sequence of p over Mab . The hypothesis implies that $p^{(\omega)}$ and q commute (as \otimes is associative). In particular, $I + I'$ is indiscernible over b . By distality, $I + a + I'$ is also b -indiscernible. This proves that $\text{tp}(a, b/\emptyset)$ is determined.

We can do the same thing adding some parameters to the base, and thus p and q are orthogonal. \square

Proposition 5.2.17. *The theory \mathcal{T} is distal if and only if any two global invariant types p and q that commute are orthogonal.*

Proof. Lemma 5.2.16 gives one implication. Conversely, assume that \mathcal{T} is not distal. Then there is a dense indiscernible sequence I , two distinct Dedekind cuts c_1 and c_2 and a and b filling them such that $a \not\perp_I b$. By Lemma 5.2.8 (strong base change), we may assume that $I \subset M$, for M a large saturated model, and $a \models \lim(c_1^-/M)$, $b \models \lim(c_2^-/M)$. Then the types $p = \lim(c_1^-/M)$ and $q = \lim(c_2^-/M)$ have the required property. \square

Consider $p, q \in S(M)$ and assume only that p is invariant. Then $p_x \otimes q_y$ is well defined, but $q_y \otimes p_x$ does not make sense in general. We show now how to define $q_y \otimes p_x$.

Lemma 5.2.18. *Let M be κ -saturated, $\kappa \geq |\mathcal{T}|^+$. Let $p, q \in S(M)$, p being A -invariant for some $|A| < \kappa$. Then there is some $B \subset M$, $|B| < \kappa$, such that $A \subseteq B$ and for $b \models q$ and any $a, a' \in M$ such that $a, a' \models p|_B$, we have $\text{tp}(a, b/A) = \text{tp}(a', b/A)$.*

Proof. Let $\mathbf{b} \models \mathbf{q}$.

We try to build inductively a sequence $\langle B_\alpha, \mathbf{a}_\alpha^0, \mathbf{a}_\alpha^1, \bar{\mathbf{c}}_\alpha, \phi_\alpha(\mathbf{x}, \mathbf{y}; \bar{\mathbf{z}}_\alpha) : \alpha < |T|^+ \rangle$ such that :

- $B_0 = A$;
- $B_\lambda = \bigcup_{\alpha < \lambda} B_\alpha$;
- $B_{\alpha+1} = B_\alpha \cup \{\mathbf{a}_\alpha^0, \mathbf{a}_\alpha^1\}$;
- $\mathbf{a}_\alpha^0, \mathbf{a}_\alpha^1 \models \mathbf{p}|_{B_\alpha}$;
- $\bar{\mathbf{c}}_\alpha \subset A$;
- $\models \phi_\alpha(\mathbf{b}, \mathbf{a}_\alpha^0; \bar{\mathbf{c}}_\alpha) \wedge \neg \phi_\alpha(\mathbf{b}, \mathbf{a}_\alpha^1; \bar{\mathbf{c}}_\alpha)$.

Assume we succeed. Then we may assume that $\phi_\alpha(\mathbf{x}, \mathbf{y}; \bar{\mathbf{z}}_\alpha) = \phi(\mathbf{x}, \mathbf{y}; \bar{\mathbf{z}})$ for all α . For $n < \omega$, let $\eta(n) = 0$ if n is even and 1 otherwise. Then the sequence $\langle \mathbf{a}'_n = \mathbf{a}_n^{\eta(n)} : n < \omega \rangle$ is indiscernible over A but $\phi(\mathbf{b}, \mathbf{a}'_n; \bar{\mathbf{c}}_n)$ holds if and only if n is even. This contradicts Proposition 5.1.10 (Shrinking of indiscernibles).

The construction must therefore stop at some B_α , and setting $B = B_\alpha$, we have the required property. \square

Let $M, \mathbf{p}, \mathbf{q}$ as in the previous lemma. Let $\mathbf{b} \models \mathbf{q}$ and $A \subset M$, $|A| < \kappa$ such that \mathbf{p} is A -invariant. Let B be given by the lemma. We define $\mathbf{q}_y \otimes \mathbf{p}_x|_A$ as $\text{tp}_{\mathbf{x}, \mathbf{y}}(\mathbf{a}, \mathbf{b}/A)$ where $\mathbf{a} \models \mathbf{p}|_B$, $\mathbf{a} \in M$. By assumption, this does not depend on the choice of \mathbf{a} . One sees easily that it does not depend on B . Finally, define $\mathbf{q}_y \otimes \mathbf{p}_x \in S(M)$ by gluing together the various $\mathbf{q}_y \otimes \mathbf{p}_x|_A$, $A \subset M$.

Notice that if \mathbf{q} was invariant to begin with, then the two definitions of $\mathbf{q}_y \otimes \mathbf{p}_x$ coincide. Note also that the associativity relation : $\mathbf{p}_x \otimes (\mathbf{q}_y \otimes \mathbf{r}_z) = (\mathbf{p}_x \otimes \mathbf{q}_y) \otimes \mathbf{r}_z$ holds in all possible cases (each product is well defined if and only if at least two of $\mathbf{p}, \mathbf{q}, \mathbf{r}$ are invariant).

In other words, associated to every such M and \mathbf{p} , there is a global M -invariant type \mathbf{p}' with the following property : for every finite $\mathbf{b} \in \mathfrak{C}$, there is a small $B \subset M$ such that every realization of $\mathbf{p}|_B$ in M satisfies $\mathbf{p}'|_{A\mathbf{b}}$. Then $\mathbf{q}_y \otimes \mathbf{p}_x$ as defined above is equal to $\mathbf{p}'_x \otimes \mathbf{q}_y$. We will call \mathbf{p}' the *inverse* of \mathbf{p} over M .

The following generalizes Lemma 5.2.16, the proof is the same, using Lemma 5.2.18 to build the Morley sequence I of \mathbf{p} inside M .

Lemma 5.2.19. *Assume T is distal. Let M be κ -saturated ($\kappa \geq |T|^+$), $\mathbf{p} \in S(M)$ be A -invariant for some A of size $< \kappa$ and $\mathbf{q} \in S(M)$ be any type. If $\mathbf{p}_x \otimes \mathbf{q}_y = \mathbf{q}_y \otimes \mathbf{p}_x$, then \mathbf{p} and \mathbf{q} are orthogonal.*

5.2.3 Generically stable measures

We prove in this section that distal theories are exactly those theories in which generically stable measures are smooth. We consider this as a justification that distality is a meaningful notion. It was proved in [62] that o-minimal theories and the p -adics have this property. This latter result will be generalized in the next section, where we prove that distality can be checked in dimension 1.

We have two tools at our disposal to link indiscernible sequences of tuples to measures. In one direction, starting with an indiscernible sequence of tuples, we can form the average measure. This construction is defined in [29], extended in [62] and recalled below. In the opposite direction, starting with a generically stable measure μ (or in fact any invariant measure), we can consider the product $\mu^{(\omega)}$ in variables x_1, x_2, \dots . We then want to realize it in some way. We do this by taking smooth extensions; see the proof of Proposition 5.2.23.

Let $I = (a_t)_{t \in [0,1]}$ be an indiscernible sequence. We can define the *average measure* μ of I as the global measure defined by $\mu(\phi(x)) = \lambda_0(\{t \in [0,1] : a_t \models \phi(x)\})$, where λ_0 is the Lebesgue measure. That measure is generically stable (in fact definable and finitely satisfiable over I).

The *support* of a measure $\mu \in \mathcal{M}(A)$ is the set of weakly-random types for μ , namely the set of types $p \in S(A)$ such that $p \vdash \neg\phi(x)$ for every formula $\phi(x) \in L(A)$ such that $\mu(\phi(x)) = 0$. We will denote it by $S(\mu)$.

Lemma 5.2.20. *Let μ be the average measure of the indiscernible sequence $I = (a_t)_{t \in [0,1]}$. Then the support $S(\mu)$ of μ is exactly the set of limit types of cuts of I .*

Proof. Let X be the set of limit types of polarized cuts of I . We first show that $S(\mu) = \bar{X}$, the closure of X . Let $\phi(x) \in L(\mathcal{C})$ have positive measure. Then by definition of μ , there is a non trivial open interval J of $[0,1]$ such that $\phi(a_t)$ holds for $t \in J$. So for any cut c in that interval, the limit type associated contains ϕ . Conversely, if $\phi(x)$ is satisfied by some $\lim(c)$, c a cut in I , then $\phi(x)$ holds on a subsequence, cofinal in c , and therefore has positive measure.

To see that X is closed, take $p \in \bar{X}$. If I is totally indiscernible, then X has one element, so assume this is not the case. Therefore the sequence I is ordered by some formula $\phi(x, y) \in L(\mathcal{C})$. For every t , p must satisfy $\phi(x, a_t) \vee \phi(a_t, x)$. We may therefore associate to p a cut $c_p = (I_1, I_2)$ of I such that p satisfies $\bigwedge_{a \in I_1} \phi(a, x) \wedge \bigwedge_{a \in I_2} \phi(x, a)$. Without loss, assume that I_1 has a maximal element a_{t_*} . For every formula $\psi(x) \in L(\mathcal{C})$ such that $\lim(c_p) \vdash \psi(x)$, there is some $t_0 > t_*$ such that for $t_* < t < t_0$, we have $a_t \models \psi(x)$. Therefore for every $q \in X$, we have $q \vdash \psi(x) \vee \phi(x, a_{t_*}) \vee \phi(a_{t_0}, x)$. It follows that p satisfies the same formula, and therefore $p \vdash \psi(x)$. We conclude that $p = \lim(c_p)$. \square

Proposition 5.2.21 (Smooth measures imply distality). *Let I be an indiscernible sequence indexed by $[0,1]$, and μ be the average measure of I . Then μ is smooth if and only if I is distal.*

Proof. Assume μ is not smooth and I is distal. Then there exists a formula $\phi(x, a)$ such that the set of $p \in S(\mathcal{C})$ such that p neither implies $\phi(x, a)$ nor its negation has positive measure (in other words, $p \in \partial\phi$). We know that the support of μ is exactly the limit types of cuts in I . Therefore, one can find ω such cuts $(c_i)_{i < \omega}$ in $\partial\phi$. Remove countably many points from I (thus not affecting any limit types) so that the cuts c_i become Dedekind.

Restricting to some sub-interval of $[0, 1]$, we may assume that $\phi(x, a)$ has constant truth value on I . Without loss, it holds on all members of I . For each index i , as $\lim(c_i) \in \partial\phi$, there is b_i filling the cut c_i over I such that $\phi(b_i, a)$ holds. As I is distal, the sequence formed by adding all the b_i to I is still indiscernible. But then the formula $\phi(x, a)$ has infinite alternation number.

Conversely, assume that I is not distal. If J is an indiscernible sequence, we write J' for the sequence J with the endpoints removed. We can find a partition $I = I_1 + I_2 + I_3$ and points b_1, b_2 such that $I'_1 + b_1 + I'_2 + I'_3$ and $I'_1 + I'_2 + b_2 + I'_3$ are indiscernible, but $I'_1 + b_1 + I'_2 + b_2 + I'_3$ is not. Without loss, assume that I_1 and I_2 have no last element. By strong base change, we may assume that the types of b_1 and b_2 over M are respectively $\lim(I_1)$ and $\lim(I_2)$. There is a formula ϕ , parameters $i_k \subset I_k$ and b'_1 realizing the same type as b_1 over M such that $\phi(i_1, b_1, i_2, b_2, i_3) \wedge \neg\phi(i_1, b'_1, i_2, b_2, i_3)$ holds. Then the border $\partial\phi$ of $\phi(i_1, x, i_2, b_2, i_3)$ contains all limit types of cuts between i_1 and i_2 and has non zero measure. This proves that μ is not smooth. \square

Corollary 5.2.22. *If all generically stable measures are smooth, then T is distal.*

We now show the converse.

Proposition 5.2.23. *If T is distal, then all generically stable measures are smooth.*

Proof. Take μ a generically stable measure over some $|T|^+$ -saturated model N . The unique global invariant extension of it will also be denoted by μ . Let a be a tuple. Let μ' be an extension of μ to Na . Take a smooth extension μ'' of μ' to some $B \supseteq Na$. Let $\langle (B_i, a_i) : i < \omega \rangle$ be a coheir sequence in $\text{tp}(B, a/N)$. Define the measures $\mu_{x_i}^i$ such that μ^i is smooth over B_i and defined over B_i the same way μ'' is over B .

Consider the measure $\lambda_{\langle x_i, i < \omega \rangle}$ defined as $\bigotimes_{i < \omega} \mu_{x_i}^i$ (this does not depend on the order of the factors since the μ^i are generically stable).

Claim : The measure $\lambda_{x_1, x_2, \dots}$ is totally indiscernible over N .

Note that $\text{tp}(B_2/B_1N)$ is non-forking over N . In particular $\mu_{x_2}^2|_{B_1N}$ does not fork over N (as it is finitely satisfiable in B_2) so by Proposition 3.3 of [29], it is the unique invariant extension of μ_{x_2} over B_1N . Therefore $\mu_{x_1}^1 \otimes \mu_{x_2}^2|_N$ is equal to $\mu_{x_1} \otimes \mu_{x_2}|_N$. Iterating, $\lambda|_N = \mu^{(\omega)}|_N$. As μ is generically stable, $\lambda_{x_1, x_2, \dots}$ is totally indiscernible over N .

Now define a measure $\eta_{(x_1, y_1), (x_2, y_2), \dots}$ over N , where y_i is a variable of the same size as B , by $\eta(\phi(x_1, x_2, \dots; y_1, y_2, \dots)) = \lambda(\phi(x_1, x_2, \dots; B_1, B_2, \dots))$. By construction, η is a measure of an indiscernible sequence. Corollary 5.2.15 works equally well with measures instead of points, with the same proof, and yields that for any increasing $\sigma : \omega \rightarrow \omega$, and any $\phi(x_1, x_2, \dots; y_1, y_2, \dots)$,

$$\eta(\phi(x_1, x_2, \dots; y_{\sigma 1}, y_{\sigma 2}, \dots)) = \eta(\phi(x_1, x_2, \dots; y_1, y_2, \dots)).$$

Therefore $\mu'|_{Na} = \mu^2|_{Na} = \mu|_{Na}$ and $\text{tp}(a/N)$ and $\mu|_N$ are orthogonal. This proves that μ is smooth. \square

We end this section by giving some type-by-type versions of Proposition 5.2.23.

Proposition 5.2.24. *Let M be a κ -saturated model ($\kappa \geq |T|^+$), $p \in S(M)$ orthogonal to all generically stable measures. Let $q \in S(M)$ be an invariant type such that $p_x \otimes q_y = q_y \otimes p_x$, then p and q are orthogonal.*

Proof. By associativity of \otimes , p and $q^{(\omega)}$ commute. Assume p and q are not orthogonal, let $(a, b) \models p \times q$ such that $\text{tp}(a, b/M) \neq p \otimes q$. Let $\phi(a, b) \in L(M)$ witness this. Without loss, ϕ has parameters in $A \subset M$, $|A| < \kappa$, q is A -invariant and for every $I, I' \subset M$, Morley sequences of q over A indexed by ω , $\text{tp}(I/Aa) = \text{tp}(I'/Aa)$. Let I_1 be a dense countable Morley sequence of q over A inside M and I_2 a dense Morley sequence of q over Mab . Then $I_1 + I_2$ is indiscernible over Aa and $I_1 + b + I_2$ is indiscernible over A . Let t be a polarized cut inside I_1 and r be the limit type of t over M .

Claim : r and p are not orthogonal, in fact, there is $c \models r$ such that $\phi(a, c)$ holds.

Assume no such c exists. Then by compactness, there is $\psi(y) \in r$ such that $p(x) \wedge \phi(x, y) \wedge \psi(y)$ is inconsistent. Let J be an interval of I_1 on which ψ holds. Let c_1 be a cut interior to J . By sliding, there is c filling c_1 (over A) such that $\phi(a, c)$ holds. Then $\neg\psi(c)$ holds. By saturation, there is $c_1 \in M$ filling c_1 over A such that $\neg\psi(c_1)$ holds. By assumption on A , $I_1 \cup \{c_1\}$ is indiscernible over Aa . We can then iterate with another cut c_2 , and after ω steps, $\psi(x)$ has infinite alternation rank, contradicting NIP, so the claim is proved.

As this holds for any cut t , if I_1 is an indiscernible segment of average μ , then $\partial\phi(a, y)$ has μ -measure one, so μ and p are not orthogonal. \square

Proposition 5.2.25. *Let p be a global invariant type, and I a Morley sequence of p . Then I is distal if and only if $p^{(\omega)}$ is orthogonal to all generically stable measures.*

Proof. The proof of Proposition 5.2.23 shows that if p is non-orthogonal to a generically stable measure, then I is not distal. (Take in the proof a realizing p and instead of taking a coheir sequence (B_i) take a non-forking indiscernible sequence (B_i) such that the corresponding sequence (a_i) is a Morley sequence of p .)

Conversely, assume p is M -invariant and not distal. Let I be a Morley segment of p over M and μ the average measure of I . Then by strong base change, $p^{(\omega)}$ and μ are not orthogonal. \square

5.2.4 Reduction to dimension 1

The goal of this section is to prove the following theorem.

Theorem 5.2.26. *If all sequences of 1-tuples are distal, then T is distal.*

We first give an informal (and incomplete) proof using measures. Assume all sequences of 1-tuples are distal and consider a generically stable measure μ . Then looking at the proof of Proposition 5.2.23 we see that μ is orthogonal to all 1-types. Then by induction, adding the points one-by-one, μ is orthogonal to every n -type. However, to

be made rigorous this proof seems to require the fact that no type forks over its base. To avoid this hypothesis and the use of measures, we give a purely combinatorial proof.

So we start with a witness of non-distality of the following form :

- a base set of parameters A , and it what follows we work over A (even when not explicitly mentioned) ;
- an indiscernible sequence $I = (a_i)_{i \in J}$ with $J = (0, 1)$ (the usual interval of \mathbb{R}) for simplicity ;
- a tuple $b = (b_j)_{j < n}$, some $l \in (0, 1)$ and tuple a such that :
 - a fills the cut “ l^+ ” : $((a_i : i \leq l), (a_i : i > l))$ of I ,
 - I is b -indiscernible,
 - I with a_l replaced by a is not indiscernible over b .

We make some simplifications. First let $m < n$ be the first integer such that $b' = b_{<m}$ satisfies the requirements in place of b . We can add $b_{<m-1}$ as parameters to the base (by base change, or equivalently we can replace a_i by $a'_i = a_i \wedge b_{<m-1}$) and replace b by b_{m-1} . Therefore, we may assume that $|b| = 1$. Next, adding again some parameters to the base, we may assume that for $i \in J$, $\text{tp}(a/b) \neq \text{tp}(a_i/b)$.

The goal of the construction that follows is to reverse the situation of a and b , *i.e.*, to construct an indiscernible sequence starting with b that is not distal, the non-distality being witnessed by a (or a conjugate of it).

Step 1 : Derived sequence

Let $r = \text{tp}(a, b)$. We construct a new sequence $(a'_i)_{i \in J}$ such that :

- a'_i fills the cut i^+ of I ;
- $\text{tp}(a'_i, b) = r$ for each i ;
- The sequence $\langle (a_i, a'_i) : i \in J \rangle$ is b -indiscernible.

This is possible by indiscernability of $(a_i)_{i \in J}$ over b (by sliding, we may choose the a'_i s filling the cuts and then extract).

Step 2 : Constructing an array

Using Lemma 5.2.8 we can iterate this construction to obtain an array $\langle a_i^n : i \in J, n < \omega \rangle$ and sequence $\langle b_n : n < \omega \rangle$ such that :

- $a_i^0 = a_i$ for each i ;
- for each $i \in J$, $0 < n < \omega$, the tuple a_i^n realizes the limit type of the cut i^+ of I over $\langle b_k, a_i^k : i \in J, k < n \rangle$;
- for each $0 < n < \omega$, $\text{tp}(b_n, (a_i^n)_{i \in J}/I) = \text{tp}(b, (a'_i)_{i \in J}/I)$.

Claim : For every $\eta : J_0 \subset J \rightarrow \omega$ injective, the sequence $\langle a_i^{\eta(i)} : i \in J_0 \rangle$ is indiscernible, of same EM-type as I .

Proof. Easy, by construction. □

Expanding and extracting, we may assume that the sequence of rows $\langle b_n + (a_i^n)_{i \in J} : 0 < n < \omega \rangle$ is indiscernible and that $\langle (a_i^n)_{0 < n < \omega} : i \in J \rangle$ is indiscernible over the sequence $(b_n)_{n < \omega}$.

Step 3 : Conclusion

Claim : The sequences $(b_n)_{n < \omega}$ and $\langle (a_i^n)_{i \in J} : 0 < n < \omega \rangle$ are weakly linked (Definition 5.2.12).

Proof. Assume for example that some $\phi(b_n, a_i^k)$ holds for all $i \in J$ and any (k, n) such that $k < n$. Take n very large and take η as in the first claim such that the truth value of “ $\eta(i) < n$ ” alternates more times than the alternation number of ϕ . Then we see that $\phi(b_n, a_i^k)$ must hold also for $k > n$ (otherwise $\phi(b_n, y)$ would alternate too much on the sequence $(a_i^{\eta(i)})$). We can do something similar if the formula ϕ has extra parameters from the b_n ’s or a_i^n ’s, thus it follows that the sequences are weakly linked. \square

Choose an increasing map $\eta : \omega \rightarrow J$, then the sequences $(b_n)_{n < \omega}$ and $(a_{\eta(n)}^n)_{n < \omega}$ are weakly linked but not mutually indiscernible. This contradicts Lemma 5.2.14 and finishes the proof of Theorem 5.2.26.

Corollary 5.2.27. *If all generically stable measures in dimension 1 are smooth, then all generically stable measures are smooth.*

This generalizes results of [62] where this was proved under additional assumptions.

Corollary 5.2.28. *If T is dp-minimal and has no generically stable type (in M), then it is distal. In particular o-minimal theories and the p -adics are distal.*

Proof. Recall from 5.2.10 that in a dp-minimal theory, any indiscernible sequence of 1-tuples is either distal or totally indiscernible. \square

Appendix : types over finite sets

Perhaps the most fascinating conjecture about NIP theories is the conjecture of uniform definability of types over finite sets. It says that for every formula $\phi(x, y) \in L$ there is some formula $\psi(x, z) \in L$ such that for any tuple a of size $|y|$ and finite set B , there is $c \in B^{|z|}$ satisfying $\models \phi(b, a) \leftrightarrow \psi(b, c)$ for any $b \in B^{|x|}$. This conjecture was stated by Laskowski and is linked with *compression schemes* arising in computational geometry and learning theory. See [34]. This conjecture was proved for dp-minimal theories by Guingona in [23].

In the case of distal theories, we state a stronger conjecture :

Conjecture 5.2.29. *Let T be distal, and let $\phi(x, y) \in L$. Then there is an integer N_ϕ such that for every finite $B \subset \mathfrak{C}$ and tuple $a \in \mathfrak{C}^{|y|}$, there is a subset $B_0 \subset B$ of size at most N_ϕ such that $\text{tp}(a/B_0) \vdash \text{tp}_\phi(a/B)$.*

Observe that the converse of the conjecture is easy. In fact, if we restrict ourselves to finite subsets B that are the underlining sets of an indiscernible sequence of tuples, then the property described in the conjecture is easily seen to be equivalent to distality of T .

The conjecture has been checked by Guingona in the case of dp-minimal linearly ordered theories (unpublished).

§5.3 Domination in non-distal theories

We have now two extreme notions for indiscernible sequences : distality and total indiscernability. We want to understand the intermediate case. This part is essentially independent of the previous one but is of course motivated by it. We first concentrate on indiscernible sequences, and then adapt the results to invariant types, where statements become simpler. A last subsection gives an application to externally definable sets.

The reader might find it useful to have in mind the example of a colored order as defined in the introduction while reading this section.

We will sometimes work with saturated indiscernible sequences, as defined below.

Definition 5.3.1 (Saturated sequence). An indiscernible sequence of α -tuples is saturated if it is indexed by an $(|T| + |\alpha|)^+$ -saturated dense linear order without end points.

In this section, all cuts are implicitly assumed to be Dedekind (*i.e.*, of infinite cofinality from both sides).

If \bar{a} fills a cut \mathfrak{c} of I , an extension $J \supseteq I$ is *compatible with \bar{a}* if \bar{a} also fills a cut of J .

We fix a global A -invariant type $\mathfrak{p} \in S_\alpha(\mathfrak{C})$, for some small parameter set A . The indiscernible sequences we will consider will be Morley sequences of \mathfrak{p} . This is not a real restriction since every indiscernible sequence is a Morley sequence of some invariant type.

The following is the main definition of this section.

Definition 5.3.2 (Domination). Let I be a dense indiscernible Morley sequence of \mathfrak{p} over A , $\mathfrak{a} \models \mathfrak{p}|_{AI}$ and \mathfrak{c} a cut of I filled by a dense sequence $\bar{a}_* = \langle \mathfrak{a}_t : t \in \mathcal{I} \rangle$ of α -tuples. We say that \bar{a}_* dominates \mathfrak{a} over (I, A) if : For every cut \mathfrak{d} of I distinct from \mathfrak{c} , and \bar{b} a dense sequence filling \mathfrak{d} , we have in the sense of $T(A)$:

$$\bar{b} \downarrow_I \bar{a}_* \Rightarrow \bar{b} \downarrow_I \mathfrak{a}.$$

We say that \bar{a}_* *strongly dominates* \mathfrak{a} over (I, A) if for every $I \subseteq J$ compatible with \bar{a}_* over A and such that $\mathfrak{a} \models \mathfrak{p}|_{AJ}$, \bar{a}_* dominates \bar{a} over J .

Notice that in this context, $\bar{b} \downarrow_I \mathfrak{a}$ means $\mathfrak{a} \models \mathfrak{p}|_{I\bar{b}}$.

EXAMPLE 5.3.3. Let T be the theory of colored orders, as defined in the introduction. Let \mathfrak{p} be an A -invariant type of an element of a new color. Let $I + \mathfrak{a}$ be a Morley sequence of \mathfrak{p} over A . Let \mathfrak{c} be a cut in I . If \mathfrak{a}_* fills \mathfrak{c} , then \mathfrak{a}_* dominates \mathfrak{a} over (I, A) if and only if \mathfrak{a} and \mathfrak{a}_* have the same color.

Lemma 5.3.4. The fact that \bar{a}_* strongly dominates \mathfrak{a} over (I, A) only depends on the similarity class of $\text{tp}(\mathfrak{a}, \bar{a}_*/I)$ over A .

Proof. The statement means that if J is a dense indiscernible sequence, \bar{b}_* and b are tuples such that $\text{tp}(b, \bar{b}_*/J)$ is similar to $\text{tp}(a, \bar{a}_*/I)$ over A , then \bar{b}_* strongly dominates b over (J, A) if and only if \bar{a}_* strongly dominates a over (I, A) . Take such \bar{b}_* , b and J . Assume that $\text{tp}(\bar{b}_*, b/J)$ is similar to $\text{tp}(\bar{a}_*, a/I)$ over A . In particular, J and I have same EM-type over A , so J is also a Morley sequence of p over A . It also follows that $b \models p|_{JA}$ so it makes sense to ask for domination.

Assume that \bar{b}_* does not strongly dominate b over (J, A) . Then we can find a dense sequence $J' \supseteq J$ compatible with \bar{b}_* such that $b \models p|_{J'A}$, some cut \mathfrak{d} of J' and sequence \bar{b}' filling \mathfrak{d} such that $\bar{b}' \perp_{J'} \bar{b}_*$, but $\bar{b}' \not\perp_{J'} b$ (all over A). By Corollary 5.1.16 (sliding), we may find $I' \supseteq I$ and \bar{a}' such that $\text{tp}(\bar{b}', \bar{b}_*, b/J')$ is similar to $\text{tp}(\bar{a}', \bar{a}_*, a/I')$ over A . This implies the following facts :

- I' is compatible with \bar{a}_* and $a \models p|_{I'A}$;
- \bar{a}' fills a cut of I' distant from the cut of \bar{a}_* ;
- $\bar{a}' \perp_{I'} \bar{a}_*$ and $\bar{a}' \not\perp_{I'} a$.

Therefore \bar{a}_* does not strongly dominate a over (I, A) . □

Lemma 5.3.5. *If \bar{a}_* strongly dominates a over (I, A) , then there is a subsequence $I' \subseteq I$ of size at most $|T| + |\alpha|$ such that \bar{a}_* strongly dominates a over (I', A) .*

Proof. This follows from the previous lemma and Lemma 5.1.11 (shrinking). □

Proposition 5.3.6. *Let I be a dense Morley sequence of p over A and $a \models p|_{AI}$, \mathfrak{c} a cut of I then there is a sequence of α -tuples \bar{a}_* of length at most $|T| + |\alpha|$ such that \bar{a}_* fills \mathfrak{c} and \bar{a}_* strongly dominates a over (I, A) .*

Proof. Recall the notation $T_I(a, \phi)$ from Section 5.1.2. If $J \subseteq J'$ are two sequences, indiscernible over A , then for any formula ϕ for which this is well defined, we have : $T_J(a, \phi) \leq T_{J'}(a, \phi)$. We will write $J \triangleleft J'$ if for some ϕ , this inequality is strict.

Let \mathfrak{J} be the class of indiscernible sequences J such that one can find dense sequences J_1 and J_2 satisfying :

- $J_1 + J + J_2$ is a Morley sequence of p over A ;
- $a \models p|_{AJ_1J_2}$.

If we have a family $(I_i)_{i < \lambda}$ of indiscernible sequences such that $I_i \subseteq I_j$ and $I_i \triangleleft I_j$ hold for all $i < j$, then taking I_λ to be $\bigcup_{i < \lambda} I_i$, we have $I_i \triangleleft I_\lambda$ for all i . Notice in addition that if each I_i belongs to \mathfrak{J} , then it is also the case for I_λ (we can find J_1 and J_2 by compactness). As the numbers $T_{J'}(a, \phi)$ are finite, it follows that we can find some sequence J in the class \mathfrak{J} such that there is no $J' \supset J$ in this class with $J \triangleleft J'$. By shrinking, we may assume that J is of size $|T| + |\alpha|$. Take J_1 and J_2 as in the definition of \mathfrak{J} . Write $\mathfrak{c} = (I_1, I_2)$. Without loss, J_1 and J_2 have same order types as I_1 and I_2 respectively. Composing by an automorphism over Aa , we may assume that $J_1 = I_1$ and $J_2 = I_2$. Then J fits in the cut \mathfrak{c} . Set $\bar{a}_* = J$.

Assume that \bar{a}_* does not strongly dominate a over (I, A) . Then there is a dense sequence $I' \supseteq I$ a cut \mathfrak{d} of I' and a sequence \bar{b} filling \mathfrak{d} such that :

- \bar{a}_* fills a cut \mathfrak{c}' of I' (over A);

- $\mathbf{a} \models \mathbf{p}|_{AI'}$;
- $\bar{\mathbf{b}} \downarrow_{I'} \bar{\mathbf{a}}_*$, and $\bar{\mathbf{b}} \not\downarrow_{I'} \mathbf{a}$.

The sequence $\mathbf{K} = I' \cup \bar{\mathbf{a}}_* \cup \bar{\mathbf{b}}$ (where $\bar{\mathbf{a}}_*$ and $\bar{\mathbf{b}}$ are placed in their respective cuts) belongs to \mathcal{J} . Also $\bar{\mathbf{b}} \not\downarrow_{I'} \mathbf{a}$ implies that $\bar{\mathbf{a}}_* \triangleleft \mathbf{K}$. This contradicts maximality of $\bar{\mathbf{a}}_*$ and proves that $\bar{\mathbf{a}}_*$ strongly dominates \mathbf{a} over (I, A) . \square

External characterization and base change

Similarly to what we did in the distal case, we give an external characterization of domination.

Proposition 5.3.7 (External characterization of domination). *Let I be a dense Morley sequence of \mathbf{p} over A , $\mathbf{a} \models \mathbf{p}|_A$. Let $\bar{\mathbf{a}}_*$ fill a cut \mathbf{c} of I over A such that $\bar{\mathbf{a}}_*$ strongly dominates \mathbf{a} over (I, A) . Let also $\mathbf{d} \in \mathcal{C}$. Assume :*

\odot *There is a partition $I = J_1 + J_2 + J_3 + J_4$ such that J_2 and J_4 are infinite, \mathbf{c} in interior to J_2 , $J_2 \cup \{\bar{\mathbf{a}}_*\}$ is indiscernible over $\text{Ad} + J_1 + J_3 + J_4$ and J_4 is a Morley sequence of \mathbf{p} over $\text{Ad} + J_1 + J_2 + J_3$.*

Then $\mathbf{a} \models \mathbf{p}|_{\text{Ad}}$.

Proof. Let $I, \mathbf{a}, \bar{\mathbf{a}}_*, \mathbf{d}, J_1, \dots, J_4$ as in the statement of the proposition. We may freely enlarge the sequence J_2 , so we may assume that it is saturated (for example, add realizations of limit types of cuts in J_2 over everything. This maintains the hypothesis).

Assume \mathbf{a} does not realize \mathbf{p} over Ad . So there is $\phi(\mathbf{d}, \bar{\mathbf{i}}; \mathbf{x}) \in L(A)$, ($\bar{\mathbf{i}} \subset I$) a formula satisfied by \mathbf{a} that witnesses it. Incorporating $\bar{\mathbf{i}}$ in \mathbf{d} and changing the partition so that $J_2 \cup J_4$ contains no point from $\bar{\mathbf{i}}$, we may assume that $\bar{\mathbf{i}} = \emptyset$. Pick a sequence of cuts of J_2 $\mathbf{c}_0 < \mathbf{c}_1 < \dots$. Let $\langle \bar{\mathbf{a}}_*^k : k < \omega \rangle$ fill the polycut $\langle \mathbf{c}_k : k < \omega \rangle$ over $\text{Ad} \cup \{J_l : l \neq 2\}$, where each $\bar{\mathbf{a}}_*^k$ is a sequence of same order type as $\bar{\mathbf{a}}_*$. Let I' denote the sequence I with the points $\bar{\mathbf{a}}_*^k$, $k > 0$, placed in their respective cuts.

Then $\text{tp}(\bar{\mathbf{a}}_*^0, \mathbf{d}/I')$ is similar to $\text{tp}(\bar{\mathbf{a}}_*, \mathbf{d}/I)$. By sliding (Corollary 5.1.15; note that our sequence is already large enough, so we do not need to increase it), we find \mathbf{a}_0 such that : $\mathbf{a}_0 \models \mathbf{p}|_{AI'}$, $\phi(\mathbf{d}; \mathbf{a}_0)$ holds and $\bar{\mathbf{a}}_*^0$ strongly dominates \mathbf{a}_0 over (I', A) .

Let K_1 realize an infinite Morley sequence of \mathbf{p} over everything considered so far. Let $I_1 = I \cup \{\bar{\mathbf{a}}_*^k : k > 1\} + K_1$ (where the tuples $\bar{\mathbf{a}}_*^k$ are placed in their respective cuts). As above, we may find $\mathbf{a}_1 \models \mathbf{p}|_{AI_1}$ such that $\bar{\mathbf{a}}_*^1$ strongly dominates \mathbf{a}_1 over (I_1, A) and $\phi(\mathbf{d}; \mathbf{a}_1)$ holds. Now as $\mathbf{a}_0 \downarrow_{I_1} \bar{\mathbf{a}}_*^1$, by the domination assumption we have $\mathbf{a}_0 \downarrow_{I_1} \mathbf{a}_1$. We iterate this construction building an indiscernible sequence $I_\omega = I + K_1 + K_2 + \dots$ and points $\langle \mathbf{a}_k : k < \omega \rangle$ filling the cuts between the K_i 's and independent over I_ω such that $\phi(\mathbf{d}; \mathbf{a}_k)$ holds for each k . As by assumption $\neg\phi(\mathbf{d}; \mathbf{x})$ holds for every $\mathbf{x} \in I_\omega$, ϕ has infinite alternation rank, contradicting NIP. \square

Proposition 5.3.8 (Base change). *Let \mathbf{p} be A invariant and $A \subset B$. If I is a dense Morley sequence of \mathbf{p} over B , $\mathbf{a} \models \mathbf{p}|_B$ and $\bar{\mathbf{a}}_*$ fills a cut of I in the sense of $T(B)$, then if $\bar{\mathbf{a}}_*$ strongly dominates \mathbf{a} over (I, A) it does so over (I, B) .*

Proof. Assume that \bar{a}_* fills a cut \mathfrak{c} of I in the sense of $T(B)$ and dominates \mathfrak{a} over (I, A) . Then let \bar{d} fill a cut \mathfrak{c}' of I over B with \mathfrak{c}' distinct from \mathfrak{c} . Assume that $\bar{d} \perp_I \bar{a}_*$ over B . Then \textcircled{C} holds with d there replaced by $\bar{d}B$. By domination over (I, A) and the previous proposition, $\mathfrak{a} \models p|I \cup \bar{d}B$. This proves that \bar{a}_* dominates \mathfrak{a} over (I, B) . This remains true if we first increase I so \bar{a}_* strongly dominates \mathfrak{a} over (I, B) . \square

5.3.1 Domination for types

We now have all we need to state domination results for types over $|T|^+$ -saturated models, instead of cuts in indiscernible sequences.

We work over a fixed κ -saturated model M . By an *invariant type* we mean here a type over M , invariant over some $A \subset M$ of size less than κ .

For the following definition, recall the construction of $p_x \otimes q_y$ when q is invariant (Lemma 5.2.18 and the paragraph following it).

Definition 5.3.9 (Distant). Let $p, q \in S(M)$ be two types, assume that at least one of them is invariant, then we say that p and q are distant if they commute : $p_x \otimes q_y = q_y \otimes p_x$ ². If $\mathfrak{a}, \mathfrak{b} \in \mathfrak{C}$, we will say that \mathfrak{a} and \mathfrak{b} are distant over M if $\text{tp}(\mathfrak{a}/M)$ and $\text{tp}(\mathfrak{b}/M)$ are.

Keep in mind that the notion “ \mathfrak{a} and \mathfrak{b} are distant over M ” only depends on $\text{tp}(\mathfrak{a}/M) \cup \text{tp}(\mathfrak{b}/M)$ and does not say anything more about $\text{tp}(\mathfrak{a}, \mathfrak{b}/M)$. In particular, in a stable theory, any \mathfrak{a} is distant from itself. So distant should not be confused with independent as defined now.

Definition 5.3.10 (Independent). Given two distant types $p, q \in S(M)$ and $\mathfrak{a} \models p$, $\mathfrak{b} \models q$ we say that \mathfrak{a} and \mathfrak{b} are independent over M if $\text{tp}(\mathfrak{a}, \mathfrak{b}/M) = p \otimes q$. We write $\mathfrak{a} \perp_M \mathfrak{b}$. This is a symmetric relation.

Definition 5.3.11 (S-domination). Let $p \in S(M)$ be any type, $\mathfrak{a} \models p$. A tuple \mathfrak{b} s-dominates \mathfrak{a} over M if :

\textcircled{S} For every invariant type $r \in S(M)$ distant from p and q , and $d \models r$, if $d \perp_M \mathfrak{b}$, then $d \perp_M \mathfrak{a}$.

The reader might be concerned by the fact that this definition depends on the choice of κ (taking a smaller κ we have less invariant types to check). However, we will see later that we get an equivalent definition if we add in \textcircled{S} the condition that r is invariant over a subset of size \aleph_0 .

EXAMPLE 5.3.12. Taking again the example of a colored order, if p and q are two invariant types (of tuples), $\bar{a} \models p$ and $\bar{b} \models q$, then \bar{b} s-dominates \bar{a} over M if and only if, for every point \mathfrak{a}_0 in $\text{range}(\bar{a})$, there is a point \mathfrak{b}_0 in $\text{range}(\bar{b}) \cup M$ of the same color.

2. Recall the definition of commuting for non-invariant given after Lemma 5.2.18

The moving-away lemma

Lemma 5.3.13. *Let $p \in S(M)$ be any type, and $a \models p$. Then there is some a_* s -dominating a over M and furthermore a_* realizes some invariant type over M .*

Proof. This is similar to Proposition 5.3.6. Start with some a_* realizing an invariant type. If it does not dominate a , there is an invariant type r distant from a_* and a over M and $b \models r|Ma_*$ such that $b \not\downarrow_M a$. Replace a_* by a_*b and iterate. By Corollary 5.1.18, this construction must stop after less than $(|T| + |\alpha|)^+$ steps. \square

For applications we will also need to show that we can find such a dominating tuple distant from any given type.

Lemma 5.3.14. *Let $I \subset M$ be a dense indiscernible sequence of α -tuples and $(I_i)_{i < \lambda}$ a family of distinct initial segments of I , with $\lambda \geq (|T| + |\alpha|)^+$. For $i < \alpha$, let $p_i = \lim(I_i/M)$. Then given a type $q \in S(M)$, there is $i < \lambda$ such that p_i is distant from q .*

Proof. Observe that the types p_i pairwise commute. Then use Corollary 5.1.18 (and the remark after it). \square

Lemma 5.3.15. *Let $p, q \in S(M)$, be types of α -tuples ($|\alpha| < \kappa$) with p invariant over some small A . Let $a \models p$. Then there is $r \in S(M)$ invariant over some B of size \aleph_0 , distant from p and q and $\bar{b} \models r$ such that $|\bar{b}| \leq |T| + |\alpha|$ and \bar{b} s -dominates a over M .*

Proof. By Proposition 5.3.6 (and Lemma 5.3.5) we can find I'_0 a dense Morley sequence of p over A of size $|T| + |\alpha|$ and \bar{a}'_* such that $a \models p|AI'_0$, \bar{a}'_* fills a cut c of I'_0 and \bar{a}'_* strongly dominates a over (I'_0, A) . Let \bar{b}' be the sequence $I'_0 \cup \bar{a}'_*$ where \bar{a}'_* is placed in its cut.

Let $I \subset M$ be a saturated Morley sequence of p over A , let c be a polarized cut of I of cofinality \aleph_0 such that $\lim(c)$ is distant from q and p (using Lemma 5.3.14). We may find some $\bar{b} \equiv_{Aa} \bar{b}'$ such that \bar{b} fills the cut c of I . Let also I_0, \bar{a}_* be such that $(\bar{b}, I_0, \bar{a}_*) \equiv (\bar{b}', I'_0, \bar{a}'_*)$. So $\bar{b} = I_0 \cup \bar{a}_*$.

Let I_∞ realize an infinite Morley sequence of p over everything. The strong base change lemma (5.2.8) works equally well if instead of considering points d_i filling the cuts c_i , we take sequences \bar{d}_i . We apply this modified version with M as set of parameters, $I + I_\infty$ as indiscernible sequence, $\bar{d}_0 = \bar{b}$ and $\bar{d}_1 = a$. We conclude that we may assume that \bar{b} is a Morley sequence of $\lim(c)$ over M .

Set $r = \text{tp}(\bar{b}/M)$ and let $B \subset M$ be of size \aleph_0 such that r is B -invariant. Note that r is a power of $\lim(c)$, so it also commutes with p and q .

Let d realize any invariant type $s \in S(M)$ distant from p and r . Assume that $d \downarrow_M \bar{b}$. Let $C \subset M$ be a subset of size $< \kappa$ such that p, s and r are invariant over C . Let $I' \subset M$ be a Morley sequence of p over C indexed by some dense order J . Then $d \wedge \bar{b}$ realizes $s \otimes r$ over CI' (indeed over M). As p is distant from both r and s , by associativity of \otimes , $p^{(J)}$ commutes with $s \otimes r$. Therefore, I' realizes $p^{(J)}$ over $Cd\bar{b}$. Similarly, \bar{b} realizes r over $CI'd$, and in particular, \bar{b} is indiscernible over $CI'd$.

Furthermore, as $I' \subset M$, \bar{b} realizes r over CI' . As r commutes with p , I' realizes $p^{(J)}$ over $C\bar{b}$, a fortiori over $A\bar{b}$. But \bar{b} is a Morley sequence of p over A . Therefore $\bar{b} + I'$ is a Morley sequence of p over A .

The hypothesis of Proposition 5.3.7 are satisfied with $J_1 = J_3 = \emptyset$, $J_2 = I_0$, $J_4 = I'$ and d there equal to Cd . We conclude that $a \models p|Cd$. As this is true for every small C , d and a are independent over M . This proves that \bar{b} s -dominates a over M . \square

Remark 5.3.16. The tuple \bar{b} constructed in the previous lemma has the following additional property :

(D) For every $d \in \mathfrak{C}$ such that $\text{tp}(d/M\bar{b})$ does not fork over M , and such that $\text{tp}(\bar{b}d/M)$ commutes with p , we have $a \downarrow_M d$.

This assumption is satisfied in particular when d is distant from a and \bar{b} , and $\bar{b} \downarrow_M d$ (although d might not realize an invariant type).

Proof. We indicate how to modify the proof above. First, we take C such that p and r are invariant over C . Next take C_1 , $C \subseteq C_1 \subset M$, such that for any $J, J' \subset M$ Morley sequences of p over C_1 indexed by ω , we have $\text{tp}(J/C\bar{b}d) = \text{tp}(J'/C\bar{b}d)$. This is possible using Lemma 5.2.18. Build I' as a Morley sequence of p over C_1 . By definition of commuting, I' is a Morley sequence of p over $C\bar{b}d$. Also because $\text{tp}(d/M\bar{b})$ does not fork over M , \bar{b} is indiscernible over Md . Finally, the proof that $\bar{b} + I'$ is a Morley sequence of p over A does not change. So as above, we may apply Proposition 5.3.7 to conclude that d and a are independent over M . \square

Corollary 5.3.17. *Let $p, q \in S(M)$ be any two types of α -tuples ($|\alpha| < \kappa$) and let $a \models p$. Then there is a_* a tuple of length $\leq |T| + |\alpha|$, distant from q over M and such that a_* s -dominates a over M . Furthermore, we may assume that $\text{tp}(a_*/M)$ is invariant over a subset of size \aleph_0 .*

Proof. By Lemma 5.3.13, there is some a_{**} s -dominating a over M and realizing some invariant type. By Lemma 5.3.15, there is a tuple a_* s -dominating a_{**} over M with the required size, whose type over M is invariant over a subset of size \aleph_0 and distant from q .

We check that a_* s -dominates a over M . Let $r \in S(M)$ be an invariant type distant from a_* and a . Let $b \models r$ with $b \downarrow_M a_*$. By Lemma 5.3.15, there is b_* s -dominating b and distant from $q = \text{tp}(a \wedge a_{**}/M)$. Furthermore assume that b_* satisfies property (D). Composing by an automorphism over Mb , we may further assume that $b_* \downarrow_M a_*$. Then as a_* s -dominates a_{**} over M , we have $b_* \downarrow_M a_{**}$ and as a_{**} s -dominates a over M , $b_* \downarrow_M a$. By property (D) this implies $b \downarrow_M a$. \square

Lemma 5.3.18 (Transitivity of s -domination). *Let $a \in \mathfrak{C}$ and let a_* s -dominate a over M . Let also a_{**} s -dominate a_* over M . Then a_{**} s -dominates a over M .*

Proof. Let $d \in \mathfrak{C}$ be distant from a and a_{**} with $d \downarrow_M a_{**}$. By Corollary 5.3.17, let d_* s -dominate d over M and distant from $a \wedge a_{**}$. Composing by an automorphism over Md , we may assume that $d_* \downarrow_M a_{**}$. Then we have $d_* \downarrow_M a_*$ and $d_* \downarrow_M a$ and finally $d \downarrow_M a$. \square

EXAMPLE 5.3.19. If $p \in S(M)$ is generically stable, and $a \models p$, then a is s -dominated by itself. In the opposite situation, if p is invariant and its Morley sequence is distal, then a is s -dominated by the empty set.

S-independence

Definition 5.3.20 (S-independence). Let p, q be any types over M , let $a \models p$ and $b \models q$. We say that a and b are s -independent over M and write $a \downarrow_M^s b$ if there is a tuple a_* realizing an invariant type, s -dominating a and distant from b such that $a_* \downarrow_M b$.

Note that if a and b are distant, then $a \downarrow_M^s b$ if and only if $a \downarrow_M b$.

Proposition 5.3.21 (Existence). Let $p, q \in S(M)$ be any two types and $a \models p$. Then there is $b \models q$ such that $a \downarrow_M^s b$.

Proof. Let a_* be s -dominating a such that a_* realizes some invariant type p_* distant from q . Take b such that $\text{tp}(a_*, b/M) = p_* \otimes q$. Then by definition $a \downarrow_M^s b$. \square

Proposition 5.3.22 (Symmetry of s -independence). S -independence is symmetric : if a and b are two tuples, then $a \downarrow_M^s b$ if and only if $b \downarrow_M^s a$ if and only if there are a_*, b_* s -dominating a and b respectively, distant from each other such that $a_* \downarrow_M b_*$.

Proof. It is enough to prove the last equivalence. To see right to left, let a_{**} s -dominate a_* and be distant from b_* and b over M . Assume also that $a_{**} \downarrow_M b_*$, then by Lemma 5.3.18, a_{**} s -dominates a over M . As it is independent from b_* over M , we have $a_{**} \downarrow_M b$ as required.

Conversely, assume that $a \downarrow_M^s b$. Let a_* be a tuple s -dominating a , realizing an invariant type over M , and distant from b such that $b \downarrow_M a_*$. We can find a tuple b'_* s -dominating b distant from a, a_* and b . As $a_* \downarrow_M b$, there is $b_* \equiv_{Mb} b'_*$ such that $a_* \downarrow_M b_*$. \square

Proposition 5.3.23 (Weight is bounded). Let $(b_i)_{i < |T|^+}$ be a sequence of tuples such that $b_i \downarrow_M^s b_{<i}$ for each i , and let $a \in \mathfrak{C}$. Then there is $i < |T|^+$ such that $a \downarrow_M^s b_i$.

Proof. By Lemma 5.3.17, we can find a family $(b_i^*)_{i < |T|^+}$ such that : For each $i < |T|^+$, b_i^* realizes an invariant type r_i distant from $q := \text{tp}(a/M)$ and $r_j, j \neq i$, b_i^* s -dominates b_i over M and $b_i^* \downarrow_M b_{<i}^*$. By Corollary 5.1.18, there is $i < |T|^+$ such that $\text{tp}(b_i^*, a/M) = r_i \otimes q$. By definition, $a \downarrow_M^s b_i$. \square

The following special case of this proposition makes no reference to s -domination.

Corollary 5.3.24. Let $q \in S(M)$ be A -invariant and, for $i < |T|^+$, let $p_i \in S(M)$ be an invariant type. Assume that p_i commutes with q , for each i . Let $(b_i) \models \bigotimes p_i$ and $a \models q$. Then there is $i < |T|^+$ such that $\text{tp}(b_i, a/N) = p_i \otimes q$.

Corollary 5.3.25. Let $a, b \in \mathfrak{C}$ such that $a \not\downarrow_M^s b$, then $\text{tp}(b/Ma)$ forks over M .

Proof. Otherwise, we could find a global M -invariant extension \tilde{p} of $\text{tp}(b/Ma)$. Take $(a_i)_{i < |T|^+}$ to be a sequence of realizations of $\text{tp}(a/M)$ with $a_0 = a$ and $a_i \downarrow_M^s a_{<i}$ for each i . By invariance, if $b_* \models \tilde{p}$ over everything, for each $i < |T|^+$, $\text{tp}(b_*, a_i/M) = \text{tp}(b_*, a/M)$ and $b_* \not\downarrow_M^s a_i$. This contradicts Proposition 5.3.23. \square

Corollary 5.3.26. *Let a and b be distant over M , then $\text{tp}(a/Mb)$ forks over M if and only if $\text{tp}(b/Ma)$ forks over M if and only if $a \not\downarrow_M b$.*

Proposition 5.3.27. *Let $p \in S(M)$ be an invariant type and $q \in S(M)$ be distant from p . Let $I = (a_i)_{i < \omega}$ be a Morley sequence of p over M and $b \models q$. Then $\lim(I/Mb) = p|_{Mb}$.*

Proof. This follows easily from Proposition 5.3.23 by making the sequence I of large cardinality. \square

EXAMPLE 5.3.28 (ACVF). *Take T to be ACVF, and M a model of T . Let $p \in S(M)$ be an invariant type of a field element. By [27], Corollary 12.14, there are definable functions f and g respectively into the residue field k and the value group Γ such that letting $p_k = f_*(p)$ and $p_\Gamma = g_*(p)$, we have :*

For any $a \models p$ and $b \in \mathfrak{C}$, $\text{tp}(a/Mb) = p|_{Mb}$ if and only if $\text{tp}(f(a)/Mb) = p_k|_{Mb}$ and $\text{tp}(g(a)/Mb) = p_\Gamma|_{Mb}$.

Take such an invariant type p and $a \models p$. Then a is s -dominated by $f(a)$ since if $b \in \mathfrak{C}$ is distant from a over M , then by distality of Γ , $\text{tp}(b/M)$ and $\text{tp}(g(a)/M)$ are orthogonal.

5.3.2 The finite-co-finite theorem and application

We prove now an analog of Proposition 5.3.23 which does not require to work over a model. We prove it by reproducing the proof of that proposition in the context of domination for indiscernible sequences.

Proposition 5.3.29. *Let A be any set of parameters and let p be some global A -invariant type. Let $a \in \mathfrak{C}$. Let I be an infinite Morley sequence of p over Aa and J be an infinite Morley sequence of p over AI . Let $\phi(x; y) \in L(A)$, then the set $\{b \in J : \models \phi(b, a)\}$ is finite or co-finite in J .*

Proof. Assume not. Then we may expand I to a saturated sequence. Without loss, the formula $\phi(x, b)$ is true for $x \in I$ and pruning J , we may assume that it is false for $x \in J$. Finally, we may expand J so that $J = \langle b_i : i < |T|^+ \rangle$.

We can find sequences $\langle \bar{b}_*^i : i < |T|^+ \rangle$ such that :

- Each \bar{b}_*^i fills some cut of I , the cuts being distinct from one another, and the \bar{b}_*^i are placed independently over I ;
- for each index i , \bar{b}_*^i strongly dominates b_i over (I, A) .

(Why? First take \bar{d}_*^0 strongly dominating b_0 over (I, A) . Let $\langle b'_i : 0 < i < |T|^+ \rangle$ be a Morley sequence of p over everything. There is an automorphism σ fixing $A \cup b_0$ sending

$\langle b'_i : 0 < i < |T|^+ \rangle$ to $\langle b_i : 0 < i < |T|^+ \rangle$. Let $\bar{b}_*^0 = \sigma(\bar{d}_*^0)$. Then take \bar{d}_*^1 strongly dominating b_1 over (I, A) with $\bar{d}_*^1 \downarrow_I \bar{b}_*^0$. And iterate.)

Let I' be the sequence I with all the \bar{b}_*^i added in their respective cuts. It is an A -indiscernible sequence. By shrinking of indiscernibles, there is $I'' \subseteq I$ obtained by removing at most $|T|$ of the tuples \bar{b}_*^i from I' such that I'' is indiscernible over Aa . Without loss, assume we have not removed the tuple \bar{b}_*^0 . Then by Proposition 5.3.7 (External characterization), $b_0 \models p|_{Aa}$. This contradicts the hypothesis. \square

Theorem 5.3.30 (Finite-co-finite theorem). *Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite. Assume that $I_1 + I_3$ is A -indiscernible and take $\phi(x; a) \in L(A)$, then the set $B = \{b \in I_2 : \models \phi(b; a)\}$ is finite or co-finite.*

Proof. This follows from the previous proposition by setting p to be the limit type of I_3^* (I_3 in reverse order). \square

Note that necessarily, B in the statement of the theorem is finite if $\neg\phi(b; a)$ holds for $b \in I_1 + I_3$ and co-finite otherwise (because you can incorporate some parts of I_1 and I_3 to I_2 , also it follows from the proof). This will be used implicitly in applications.

Corollary 5.3.31. *Let $I = I_1 + I_2 + I_3$ be indiscernible, I_1 and I_3 being infinite with no endpoints and I_2 densely ordered. Assume that $I_1 + I_3$ is A -indiscernible. Write $I_2 = \langle a_i \rangle_{i \in \mathcal{J}}$. Then given some linear order $\mathcal{J} \supseteq \mathcal{J}$, one can find tuples a_i , $i \in \mathcal{J} \setminus \mathcal{J}$ such that :*

- $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{J} \rangle + I_3$ is indiscernible over A ,
- $I_1 + \langle a_i : i \in \mathcal{J} \rangle + I_3$ is indiscernible.

Proof. We construct the points a_i , $i \in \mathcal{J} \setminus \mathcal{J}$ simply by realizing limit types of cuts of I_2 over everything. More precisely, given c a cut of \mathcal{J} , identify c with the corresponding cut of I_2 . Assume for simplicity that c has infinite cofinality from the right and let p_c be $\lim(c^+)$ (seen a global type). Note that if $c \neq c'$, then the types p_c and $p_{c'}$ commute. Let \mathcal{J}_c be the convex subset of \mathcal{J} formed by elements falling in the cut c . Finally take $\langle a_i : i \in \mathcal{J} \setminus \mathcal{J} \rangle$ to realize $\bigotimes_c p_c^{(\mathcal{J}_c)}$ over IA .

The second condition is obviously satisfied, so we have to check the first one. We start by considering a cut c , and show that $I_1 + \langle a_i : i \in \mathcal{J}_c \rangle + I_3$ is indiscernible over A . The fact that for $i \in \mathcal{J}_c$, and $a \in I_1$, $\text{tp}(a_i/A) = \text{tp}(a/A)$ follows immediately from the finite-co-finite theorem 5.3.30. Now consider $i < j \in \mathcal{J}_c$ and $\phi(x_1, x_2) \in L(A)$ a formula. Assume that for $a \in I_1$, $b \in I_3$ we have $\models \phi(a, b)$. Write $c = (\mathcal{K}_1, \mathcal{K}_2)$, seen as a cut of \mathcal{J} . By construction of $\langle a_i \rangle_{i \in \mathcal{J}_c}$ and shrinking of indiscernibles (Proposition 5.1.10), we have :

$$\models \phi(a_i, a_j) \iff \text{for some coinital } \mathcal{K} \subset \mathcal{K}_2, \phi(a_s, a_t) \text{ holds for } s < t \in \mathcal{K}.$$

Assume we have $\neg\phi(a_i, a_j)$. So easily, we can find points $s_1 < t_1 < s_2 < t_2 < \dots \in \mathcal{K}_2$ such that $\neg\phi(a_{s_k}, a_{t_k})$ holds for each $k < \omega$. Let $L_2 = \langle a_{s_k} \wedge a_{t_k} : k < \omega \rangle$. Take also L_1 to be any sequence of increasing pairs of members of I_1 , so that $L_1 + L_2$ is indiscernible, and pick similarly L_3 . Then the finite-co-finite theorem applied to the sequence $L_1 + L_2 + L_3$ gives us a contradiction.

We can do the same reasoning if $\phi(x_1, x_2)$ has parameters in AI_1I_2 (by adding parts of I_1I_2 to A and decreasing them). Also one sees at once that the construction generalizes to formulas $\phi(x_1, \dots, x_n)$ with more variables and we obtain that $I_1 + \langle a_i : i \in \mathcal{J}_c \rangle + I_3$ is indiscernible over A .

Next, we look at two cuts $c_1 < c_2$ and we want to see that $I_1 + \langle a_i : i \in \mathcal{J}_{c_1} + \mathcal{J}_{c_2} \rangle + I_3$ is indiscernible over A . We know that $\langle a_i : i \in \mathcal{J}_{c_2} \rangle$ realizes $p_{c_2}^{(\mathcal{J}_{c_2})}$ over everything else. We may assume that \mathcal{J}_{c_1} is without endpoints. Take some finite $\mathcal{K}_0 \subset \mathcal{J}_{c_1}$ and let \mathcal{K}_1 be $\{i \in \mathcal{J}_{c_1} : i > \mathcal{K}_0\}$. Then the sequence $\langle a_i : i \in \mathcal{K}_1 \rangle + I_3$ is indiscernible over $A \cup \{a_i : i \in \mathcal{K}_0\}$. The same reasoning as above shows that the sequence $\langle a_i : i \in \mathcal{K}_1 \rangle + \langle a_i : i \in \mathcal{J}_{c_2} \rangle + I_3$ is indiscernible over $A \cup \{a_i : i \in \mathcal{K}_0\}$. It follows that $I_1 + \langle a_i : i \in \mathcal{J}_{c_1} + \mathcal{J}_{c_2} \rangle + I_3$ is indiscernible over A .

Iteratively, we prove that $I_1 + \langle a_i : i \in \mathcal{J}_{c_1} + \dots + \mathcal{J}_{c_n} \rangle + I_3$ is indiscernible over A and finally, that $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{J} \rangle + I_3$ is indiscernible over A . \square

Corollary 5.3.32. *Let $I_1 + I_2 + I_3$ be an indiscernible sequence of finite tuples, with I_1 and I_3 infinite without endpoints. Assume that $I_1 + I_3$ is indiscernible over A . Then we can find some subsequence $I'_2 \subset I_2$ with $I_2 \setminus I'_2$ of size at most $|T| + |A|$ such that $I_1 + I'_2 + I_3$ is indiscernible over A .*

Proof. Without loss, we may assume that I_2 is densely ordered. Write $I_2 = \langle a_i : i \in \mathcal{J} \rangle$ and take some $|\mathcal{J}|^+$ -saturated linear order $\mathcal{J} \supset \mathcal{J}$. By Corollary 5.3.31 we can find tuples $\langle a_i : i \in \mathcal{J} \setminus \mathcal{J} \rangle$ such that :

- $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{J} \rangle + I_3$ is indiscernible over A ,
- $I_1 + \langle a_i : i \in \mathcal{J} \rangle + I_3$ is indiscernible.

By shrinking of indiscernibles, there is $\mathcal{J}_0 \subset \mathcal{J}$ of size at most $|T| + |A|$ such that $I_1 + \langle a_i : i \in \mathcal{J} \setminus \mathcal{J}_0 \rangle + I_3$ is indiscernible. Then set $I'_2 = \langle a_i : i \in \mathcal{J} \setminus \mathcal{J}_0 \rangle$. \square

We now give an application of this result to externally definable sets.

We will use the following notation : if $M \models T$, $M \prec N$ is an elementary extension and $A \subseteq N$ containing M , then $M_{[A]}$ is the structure with universe M with language composed of a predicate for every subset of M^l (any l) of the form $\phi(M; \bar{c})$, $\bar{c} \in A^k$ for any $\phi(\bar{x}; \bar{y}) \in L(M)$, interpreted the obvious way.

Shelah proved in [56] that $M_{[\mathcal{C}]}$ eliminates quantifiers. We refer the reader to [16] for a slightly different approach, that we will use (and recall) here. If $p \in S(M)$ is any type and $a \models p$, then it is not true in general that $M_{[a]}$ eliminates quantifiers (see [16], Example 1.8 for a counterexample). However it is conjectured in [16] that $M_{[I]}$ does, where I is a coheir sequence starting with a . We prove a special case of this when p is interior to M . See the definition below.

We will need some notions from [16] that we recall now. If X is an externally definable subset of X (i.e., a subset of the form $\phi(M, c)$ for some tuple $c \in \mathcal{C}$), then an *honest definition* of X is a formula $\theta(x, d) \in L(\mathcal{C})$ such that (1) $\theta(M, d) = X$ and (2) for every formula $\psi(x) \in L(M)$ such that $X \subseteq \psi(M)$ then $\mathcal{C} \models \theta(x) \rightarrow \psi(x)$.

Lemma 5.3.33. *If $A \subset \mathcal{C}$ containing M is such that for every formula $\phi(x; c) \in L(A)$, $\phi(M; c)$ has an honest definition with parameters in A , then $M_{[A]}$ eliminates quantifiers.*

Proof. Let $\phi(x, y; c) \in L(A)$ and let $\theta(x, y; d) \in L(A)$ be an honest definition of $X := \phi(M; c)$. Let $\pi : M^{|x|+|y|} \rightarrow M$ be the projection on the first $|x|$ coordinates. Let $\psi(x; d) = \exists y \theta(x, y; d)$. Then $\psi(M; d) = \pi(X)$: it is clear that $\psi(M; d) \subseteq \pi(X)$, and if $a \in M^{|x|} \setminus \pi(X)$, then the set $\{(x, y) \in M^{|x|+|y|} : y \neq a\}$ contains X and by honesty $\mathfrak{C} \models \theta(x, y; d) \rightarrow y \neq a$ which gives the reverse inclusion. \square

Definition 5.3.34. Let p be an M -invariant global type. We say that p is *interior* to M if $p^{(\omega)}$ is both an heir and a co-heir of its restriction to M .

An example of an interior type is given by the following situation : let $I \subset M$ be indiscernible and c a cut interior to I such that M respects c . Then the type $p = \lim(c^+)$ is interior to M .

Lemma 5.3.35. *Let p be a global M -invariant type interior to M . Let $I_0 + I_1 + I_2$ be a Morley sequence of p over M . For $i < 3$ let $\bar{a}_i \subset I_i$ be a finite tuple. Assume that $\bar{a}_1 \models \phi(\bar{x}; \bar{a}_0, \bar{a}_2)$, $\phi \in L(M)$, then there are two tuples $\bar{b}_0, \bar{b}_2 \subset M$ such that $\bar{a}_1 \models \phi(\bar{x}; \bar{b}_0, \bar{b}_2)$.*

Proof. First find \bar{b}_2 such that $\bar{a}_1 \models \phi(\bar{x}; \bar{a}_0, \bar{b}_2)$ by the coheir hypothesis. Then find \bar{b}_0 by the heir hypothesis. \square

Theorem 5.3.36 (Shelah expansion for interior types). *Let p be a global M -invariant type interior to M . Let I be a Morley sequence of p over M . Then $M_{[I]}$ eliminates quantifiers.*

Proof. Take a saturated extension $M_{[I]} \prec N^*$ of size $\kappa > |M|$. The model N^* can be seen as a reduct to the language of $M_{[I]}$ of some $N_{[J]}$ for $M \prec N$ and $J \equiv_M I$, J indiscernible over N . Without loss $I = J$. Notice that N^* and $N_{[I]}$ have the same definable sets.

Claim : There is an indiscernible sequence $I_1 + I_2 \subset N$ such that N respects the cut $c = (I_1, I_2)$ and $I \models \lim(c^+)^{(\omega)}$.

Proof : Write $N = \bigcup_{i < \kappa} A_i$ with $|A_i| < \kappa$. Let $i < \kappa$. By Lemma 5.3.35 and saturation, we can find sequences $K_i, L_i \subset N$ of order type ω such that $K_i + I + L_i$ is indiscernible over A_i . Let $I_1 = K_1 + K_2 + \dots$ and $I_2 = \dots + L_2 + L_1$, the sums ranging over $i < \kappa$. The required property is then easy to check.

Let $\phi(x; y)$ be a formula and $a_0 \models p$, $a_0 \in I$. We consider the pair (M, N) and show that $\phi(a_0; M)$ has an honest definition with parameters in $M + I_1 + I_2$.

By the Theorem 5.3.30 and compactness, there are integers n, N and a finite set of formulas δ such that for every finite sequence $J_1 + J_3 + J_2$, satisfying :

- J_1 and J_2 are of size at least n ,
 - $J_1 + J_3 + J_2$ is indiscernible,
 - $J_1 + J_2$ is δ -indiscernible over b and
 - $\phi(x; b)$ holds on all elements of J_1 and J_2 ,
- then $\neg\phi(x; b)$ holds on at most N elements of J_3 .

Let $I'_1 \subset I_1$ and $I'_2 \subset I_2$ be finite of size n such that $I'_1 + I'_2$ is M -indiscernible. Consider the formula $\theta(y) \in L(MI)$ such that if $b \models \theta(y)$, then $I'_1 + I'_2$ is δ -indiscernible

over \mathbf{a} , and $\phi(\bar{\mathbf{a}}_0; \mathbf{y})$ holds on all elements of $I'_1 + I'_2$. Define analogously $\theta_1(\mathbf{y})$ using $\neg\phi$ instead of ϕ .

Then, for every $\mathbf{b} \in M$, $\theta(\mathbf{b})$ holds if and only if $\phi(\mathbf{a}_0; \mathbf{b})$ holds. Also, if $\mathbf{b} \in N$, and $\theta(\mathbf{b})$ holds, then $\phi(\mathbf{a}_0; \mathbf{b})$ holds (Why? Only finitely many elements \mathbf{a} from $I_1 + I_2$, with $I'_1 < \mathbf{a} < I'_2$ can satisfy $\phi(\mathbf{a}; \mathbf{b})$). This easily implies that θ is an honest definition of $\phi(\mathbf{a}_0; M)$.

To conclude the theorem, notice that we can do the same thing replacing \mathbf{p} by $\mathbf{p}^{(n)}$ for any n , which takes care of formulas $\phi(\bar{\mathbf{a}}; \mathbf{y})$ with $\bar{\mathbf{a}}$ a finite subset of I instead of one element. \square

As a consequence of the proof, we obtain a uniformity result for honest definitions.

Porism 5.3.37. *Let $\phi(x; \mathbf{y})$ be a formula, then there is a formula $\theta(x; \mathbf{z})$ such that : For every model M , global M -invariant type \mathbf{p} interior to M and $\mathbf{a} \models \mathbf{p}|_M$, there is some $\mathbf{c} \in \mathfrak{C}$ such that $\theta(x; \mathbf{c})$ is an honest definition of $\phi(M; \mathbf{a})$.*

§5.4 Sharp theories

In this last section, we study theories in which the ‘stable part’ of types is controlled by generically stable types. We give a definition, a criterion using indiscernible sequences and show that it is enough to check that criterion in dimension 1. One could probably introduce stronger notions, and ask for example that types are s -dominates by types living in a stable sort, but we do not pursue this here.

Definition 5.4.1. The theory T is sharp if for every $|T|^+$ -saturated model M and $\mathbf{p} \in S(M)$ an invariant type realized by \mathbf{a} , there is some generically stable type $\mathbf{q} \in S(M)$ and $\mathbf{a}_* \models \mathbf{q}$ such that \mathbf{a}_* s -dominates \mathbf{a} over M .

Definition 5.4.2. Let $I = \langle \mathbf{a}_i : i \in \mathcal{I} \rangle$ be a dense indiscernible sequence. A decomposition of I is an indiscernible sequence $K = \langle \mathbf{a}_i \wedge \mathbf{b}_i : i \in \mathcal{I} \rangle$ where the sequence $J = \langle \mathbf{b}_i : i \in \mathcal{I} \rangle$ is totally indiscernible and such that :

For every K' of same EM-type as K , \mathbf{c} a Dedekind cut of K' , $\mathbf{d} \in \mathfrak{C}$ such that K' is indiscernible over \mathbf{d} and $\mathbf{a} \wedge \mathbf{b}$ filling \mathbf{c} ; if there is \mathbf{a}' such that $\mathbf{a}' \wedge \mathbf{b}$ fills \mathbf{c} over $\mathbf{d}K'$, then $\mathbf{a} \wedge \mathbf{b}$ fills \mathbf{c} over $\mathbf{d}K$.

By usual sliding argument, if K is dense and contains some Dedekind cut \mathbf{c} , it is enough to check the condition for $K' = K$.

An indiscernible sequence $I = \langle \mathbf{a}_i : i \in \mathcal{I} \rangle$ is *decomposable* if it admits a decomposition $K = \langle \mathbf{a}_i \wedge \mathbf{b}_i : i \in \mathcal{I} \rangle$. In this case, we will say that I is decomposable over $\langle \mathbf{b}_i : i \in \mathcal{I} \rangle$.

Remark 5.4.3. There are two trivial cases of decomposability : If I is distal, then it is decomposable over the sequence of empty tuples, if I is totally indiscernible, it is decomposable over itself.

Lemma 5.4.4 (Internal characterization). *An indiscernible sequence $I = (a_i \wedge b_i)_{i \in J}$ is a decomposition, if and only if the following holds :*

☒ *For every J, K, L infinite indiscernible sequences without endpoints of same EM-type as I and $a \wedge b, a' \wedge b'$, if $J + a \wedge b + K + L, J + K + a' \wedge b' + L$ are indiscernible, and there exist a_0, a'_0 such that $J + a_0 \wedge b + K + a'_0 \wedge b' + L$ is indiscernible, then $J + a \wedge b + K + a' \wedge b' + L$ is indiscernible.*

Proof. Assume that I is a decomposition. Then taking $d = a'_0 \wedge b' + L$ in the definition, we see that $J + a \wedge b + K$ is indiscernible over $a'_0 \wedge b' + L$. Then taking $d = J + a \wedge b$, we get that $K + a' \wedge b' + L$ is indiscernible over $J + a \wedge b$, so $J + a \wedge b + K + a' \wedge b' + L$ is indiscernible.

Conversely, assume ☒ holds and without loss J is a dense order. Notice that the analog of ☒ where we fill n cuts instead of 2 follows from ☒ by easy induction (as in Lemma 5.2.6). Let $d \in \mathfrak{C}, c, a \wedge b$ and a' be as in the definition of decomposition. Assume that $a \wedge b$ does not fill c over Ad . Adding parameters to d if necessary, we may assume that for some formula $\phi(x, y)$, and all $a_* \wedge b_* \in I$, we have $\phi(a_* \wedge b_*, d) \wedge \neg \phi(a \wedge b, d)$. Fix some increasing sequence $(c_k)_{k < \omega}$ of Dedekind cuts of I . For each $k < \omega$, we can find a_k, a'_k, b_k such that $\text{tp}(a_k, a'_k, b_k, d/I)$ is similar to $\text{tp}(a, a', b, d/I)$ and $a_k \wedge b_k$ fills the cut c_k . By ☒ and the remark above, the sequence obtained by adding all the tuples $a_k \wedge b_k$ to I in their respective cuts is indiscernible. Then the formula $\phi(x, y)$ has infinite alternation rank. \square

We will need the following strengthening of Lemma 5.2.8.

Lemma 5.4.5 (Strong base change 2). *Let $I = (a_i \wedge b_i)_{i \in J}$ be an indiscernible sequence and $A \supset I$ a set of parameters. Let $(c_i)_{i \in J}$ be a sequence of pairwise distinct polarized Dedekind cuts in I . Call c'_i the corresponding cut in the sequence $(b_i)_{i \in J}$. For each i let $d_i \wedge e_i$ fill the cut c_i . Assume also that the sequence $(e_i)_{i \in J}$ realizes $\bigotimes \lim(c'_i)$ over I . Then there exist $(d'_i \wedge e'_i)_{i \in J}$ such that*

- $\text{tp}(\langle d'_i \wedge e'_i : i \in J \rangle / I) = \text{tp}(\langle d_i \wedge e_i : i \in J \rangle / I)$,
- for each i , $\text{tp}(d'_i \wedge e'_i / A) = \lim(c_i / A)$,
- $(e'_i)_{i \in J}$ realizes $\bigotimes_i \lim(c'_i)$ over A .

Proof. The proof is essentially the same as that of Lemma 5.2.8.

Assume the result does not hold. Then by compactness, we may assume that $J = \{1, \dots, n\}$ and that there is a formula $\phi(x_1 \wedge y_1, \dots, x_n \wedge y_n) \in \text{tp}(\langle d_i \wedge e_i : i \rangle / I)$, a formula $\theta(y_1, \dots, y_n) \in \bigotimes \lim(t'_i / m)$ and formulas $\psi_i(x_i, y_i) \in \lim(c_i / m)$ for some finite $m \in A$ such that $\phi(x_1 \wedge y_1, \dots, x_n \wedge y_n) \wedge \theta(y_1, \dots, y_n) \wedge \bigwedge_i \psi_i(x_i, y_i)$ is inconsistent. Let I_0 denote the parameters of ϕ .

Assume for simplicity that $n = 2$ (the proof for $n > 2$ is the same) and without loss c_i is polarized as c_i^- . For $i = 1, 2$ take $(I_i, I'_i) \trianglelefteq c_i$ such that ψ_i holds on all elements of I_i , $\theta(y_1, y_2)$ holds for each $(x_1 \wedge y_1, x_2 \wedge y_2) \in I_1 \times I_2$, and $I_i \cup I'_i$ contains no element of I_0 . Then $I_1 + I'_1$ and $I_2 + I'_2$ are mutually indiscernible over I_0 . So for every two cuts d_1 and d_2 respectively from $I_1 + I'_1$ and $I_2 + I'_2$, we can find points $d_1 \wedge e_1$ and $d_2 \wedge e_2$ filling those cuts (even seen as cuts of I) such that $\phi(d_1, e_1, d_2, e_2)$ holds and there are d'_1, d'_2 such that $(d'_1 \wedge e_1, d'_2 \wedge e_2)$ fills the polycut (d_1, d_2) over I .

Take a cut \mathfrak{d}_1 inside I_1 and \mathfrak{d}_2 inside I_2 and see them as cuts of I . Fill \mathfrak{d}_1 by $d_1 \wedge e_1$ and \mathfrak{d}_2 by $d_2 \wedge e_2$ as above. By hypothesis, either $\neg\theta(e_1, e_2)$, $\neg\psi_1(d_1, e_1)$ or $\neg\psi_2(d_2, e_2)$ holds. In one of the latter two cases, proceed as in Lemma 5.2.8. In the first case, keep e_1 and e_2 and add points (d'_1, d'_2) such that I with $d'_1 \wedge e_1$ and $d'_2 \wedge e_2$ added is indiscernible. Then iterate with $I \cup \{d'_1 \wedge e_1, d'_2 \wedge e_2\}$ instead of I .

After iterating this ω times, either ψ_1 , ψ_2 or θ has infinite alternation rank. \square

Lemma 5.4.6 (Base change). *The notion of being a decomposition is stable both ways under base change : If $(a_i \wedge b_i)_{i \in J}$ is A -indiscernible, then it is a decomposition in T if and only if it is a decomposition in $T(A)$.*

Proof. Assume $I = (a_i \wedge b_i)_{i \in J}$ is a decomposition, then it follows immediately from the definition that it is a decomposition from the point of view of $T(A)$.

For the converse, use the internal characterization and strong base change 2 (Lemma 5.4.5) as in the proof of Corollary 5.2.9. \square

Lemma 5.4.7. *Assume that $K = (a_i \wedge b_i)_{i \in J}$ is a decomposition of $I = (a_i)_{i \in J}$. Let \mathfrak{c} be a cut of K filled by a sequence L and denote by \mathfrak{c}' the corresponding cut in $(b_i)_{i \in J}$. Let L_2 the projection on L on the second factor (so L_2 is a totally indiscernible sequence). Let $\mathfrak{d} \in \mathfrak{C}$ be such that K is indiscernible over \mathfrak{d} and L_2 is a Morley sequence of the limit type of \mathfrak{c}' over $K\mathfrak{d}$. Then $K \cup L$ is indiscernible over \mathfrak{d} (where L is placed in the cut \mathfrak{c}).*

Proof. Assume L is dense of size $|T|$ and using Corollary 5.3.31, increase L to some saturated sequence L' filling \mathfrak{c} and such that the sequence $K_0 = K \cup (L' \setminus L)$ is indiscernible over \mathfrak{d} . Let now $a_1 \wedge b_1 \in L$. This tuple fills a Dedekind cut of K_0 . By domination in the sequence K_0 , we see that $K_1 = K_0 \cup \{a_1 \wedge b_1\}$ is indiscernible over \mathfrak{d} . Then we can take some other $a_2 \wedge b_2 \in L$. It fills a Dedekind cut of K_1 and by domination in K_1 , $K_2 = K_1 \cup \{a_1 \wedge b_1\}$ is indiscernible over \mathfrak{d} . Iterating, we see that $K \cup L'$ is indiscernible over \mathfrak{d} and therefore $K \cup L$ is indiscernible over \mathfrak{d} . \square

Lemma 5.4.8. *Let M be a $|T|^+$ -saturated model and $p, q \in S(M)$ be two commuting invariant types. Take $I \models p^{(\omega)}$ and any $\mathfrak{b} \models q$. Then we may find two sequences I_1, I_2 such that $I_1 + I + I_2$ is a Morley sequence of p over M and $I_1 + I_2$ is a Morley sequence of p over $M\mathfrak{b}$.*

Proof. Let r be the inverse of p over M (recall the definition as stated after Lemma 5.2.18). We take I_2 to be a Morley sequence of p over $M\mathfrak{b}$ and then I_1 to be a Morley sequence of r , indexed in the opposite order, over $M I_2 \mathfrak{b}$. Over M , the Morley sequence of r is the opposite of the Morley sequence of p so the first statement follows. To see the second statement, recall that if $s \in S(M)$ is any invariant type, then $r_x \otimes s_y|_M = s_y \otimes p_x|_M$. In particular,

$$r_x \otimes (q_y \otimes p_{x_1, \dots, x_n}^{(n)})|_M = (q_y \otimes p_{x_1, \dots, x_n}^{(n)}) \otimes p_x|_M = q_y \otimes p_{x, x_1, \dots, x_n}^{(n+1)}|_M.$$

The result follows. \square

Proposition 5.4.9. *Assume that all sequences are decomposable, then T is sharp.*

Proof. Let M be $|T|^+$ -saturated and $p \in S(M)$ be an A -invariant type. Let $a \models p$. Let $I \subset M$ be a small dense Morley sequence of p over A and let $K \subset M$ be a decomposition of I . Let c be a Dedekind cut of K and c_1 the corresponding cut of I . Construct some dense sequence L realizing a power of $\lim(c_1^+/M)$ as in the proof of the moving away lemma 5.3.15 (i.e., there is a convex $L' \subset L$ such that L' strongly dominates a over (L, A)). Extend L to V realizing a power of $\lim(c^+/M)$. So V is the union of L and some totally indiscernible sequence W . The type of W over M is generically stable.

Claim : W s -dominates a over M .

Proof : Let $d \in \mathfrak{C}$ be distant from a and independent from W over M . Let d_* realize an invariant type distant from aV over M such that d_* s -dominates d and is independent from W over M . If we show that $d_* \downarrow_M L$, then as L s -dominates a it will follow that $d_* \downarrow_M a$ and therefore $d \downarrow_M a$. Replacing d by d_* , we may now assume that d is distant from aV over M and realizes an invariant type.

Call $r = \lim(c^+)$ (a global invariant type). By Lemma 5.4.8, let I_1 and I_2 be two sequences such that $I_1 + I_2$ is indiscernible over dM and $I_1 + V + I_2$ is indiscernible over M . Also as d is independent from W over M , the hypothesis of Lemma 5.4.7 are satisfied (where L_2 there is W here). We conclude that d is independent from L over M and therefore d is independent from a over M . \square

5.4.1 Reduction to dimension 1

We prove here the following proposition.

Proposition 5.4.10. *Assume that all sequences of 1-tuples are decomposable, then every sequence is decomposable.*

Assume from now on that all indiscernible sequences of 1-tuples are decomposable. We will take an arbitrary indiscernible sequence and build a decomposition for it adjoining totally indiscernible sequences to it one-by-one. The proof is an adaptation of the one from Section 5.2.4. We start with a base set of parameters A that we allow to grow freely during the construction. In what follows, we work over A , even when not explicitly mentioned. We have an indiscernible sequence $I = \langle a_i \wedge \alpha_i : i \in \mathcal{I} \rangle$, where $\mathcal{I} = (0, 1)$ for simplicity and such that the sequence $\langle \alpha_i : i \in \mathcal{I} \rangle$ is totally indiscernible.

For every $i \in \mathcal{I}$, call c_i the cut “ i^+ ” of I and c'_i the associated cut in the sequence $\langle \alpha_i : i \in \mathcal{I} \rangle$.

Step 1 : Derived sequence

Assume we have a witness of non-decomposition in the following form :

- A tuple $b \in \mathfrak{C}$, some $j \in (0, 1)$ and a pair (a, α) such that :
- $a \wedge \alpha$ fills the cut c_j of I ,
- I is b -indiscernible,
- α realizes the type $\lim(c'_j)$ over Ib ,
- I with $a_j \wedge \alpha_j$ replaced by $a \wedge \alpha$ is not indiscernible over b .

As in Section 5.2.4, adding parameters to the base, we may assume that \mathbf{b} is a single point, and that $\text{tp}(\mathbf{a} \wedge \alpha / \mathbf{b}) \neq \text{tp}(\mathbf{a}_j \wedge \alpha_j / \mathbf{b})$. Let $\mathbf{r} = \text{tp}(\mathbf{a} \wedge \alpha, \mathbf{b})$.

We construct a new sequence $\langle \mathbf{a}'_i \wedge \alpha'_i : i \in \mathcal{J} \rangle$ such that :

- $\mathbf{a}'_i \wedge \alpha'_i$ fills the cut \mathbf{c}_i of \mathbf{I} ,
- $\text{tp}(\mathbf{a}'_i \wedge \alpha'_i, \mathbf{b}) = \mathbf{r}$ for each i ,
- The sequence $(\alpha'_i)_{i \in \mathcal{J}}$ realizes $\bigotimes_{i \in \mathcal{J}} \lim(\mathbf{c}'_i)$ over \mathbf{Ib} ,
- The sequence $\langle \mathbf{a}_i \wedge \alpha_i \wedge \mathbf{a}'_i \wedge \alpha'_i : i \in \mathcal{J} \rangle$ is \mathbf{b} -indiscernible.

This is possible by indiscernibility of $(\mathbf{a}_i \wedge \alpha_i)_i$ over \mathbf{b} (first pick the points α'_i then choose the \mathbf{a}_i filling the cuts and then extract).

Step 2 : Constructing an array

Using Lemma 5.4.5, iterate this construction to obtain an array $\langle \mathbf{a}^n_i \wedge \alpha^n_i : i \in \mathcal{J}, n < \omega \rangle$ and sequence $\langle \mathbf{b}_n : n < \omega \rangle$ such that :

- $\mathbf{a}^0_i \wedge \alpha^0_i = \mathbf{a}_i \wedge \alpha_i$ for each i ,
- For each $i \in \mathcal{J}$, $0 < n < \omega$, the tuple $\mathbf{a}^n_i \wedge \alpha^n_i$ realizes $\lim(\mathbf{c}_i)$ over $\langle \mathbf{b}_k, \mathbf{a}^k_i \wedge \alpha^k_i : i \in \mathcal{J}, k < n \rangle$,
- For each $0 < n < \omega$, the sequence $(\alpha^n_i)_{i \in \mathcal{J}}$ realizes the type $\bigotimes_{i \in \mathcal{J}} \lim(\mathbf{c}'_i)$ over $\langle \mathbf{b}_k, \mathbf{a}^k_i \wedge \alpha^k_i : i \in \mathcal{J}, k < n \rangle$,
- For each $0 < n < \omega$, $\text{tp}(\mathbf{b}_n, \langle \mathbf{a}^n_i \wedge \alpha^n_i : i \in \mathcal{J} \rangle / \mathbf{I}) = \text{tp}(\mathbf{b}, \langle \mathbf{a}'_i \wedge \alpha'_i : i \in \mathcal{J} \rangle / \mathbf{I})$.

Claim : For every $\eta : \mathcal{J}_0 \subset \mathcal{J} \rightarrow \omega$ injective, the sequence $\langle \mathbf{a}^{\eta(i)}_i \wedge \alpha^{\eta(i)}_i : i \in \mathcal{J}_0 \rangle$ is indiscernible, of same EM-type as \mathbf{I} .

The sequence $\mathbf{U} = \langle \alpha^n_i : (i, n) \in \mathcal{J} \times \omega \rangle$, where $\mathcal{J} \times \omega$ is ordered lexicographically, is totally indiscernible.

Proof. Easy, by construction. □

Expanding and extracting, we may assume that the sequence of rows $\langle \mathbf{b}_n + (\mathbf{a}^n_i \wedge \alpha^n_i)_{i \in \mathcal{J}} : 0 < n < \omega \rangle$ is indiscernible. By assumption all sequences of points are decomposable. So let $(\mathbf{b}_n \wedge \beta_n)_{n < \omega}$ be an decomposition of $(\mathbf{b}_n)_{n < \omega}$. Expanding and extracting again, we may assume that the new sequence of rows $\langle \mathbf{b}_n \wedge \beta_n + (\mathbf{a}^n_i \wedge \alpha^n_i)_{i \in \mathcal{J}} : 0 < n < \omega \rangle$ is indiscernible and that the sequence of columns $\langle (\mathbf{a}^n_i \wedge \alpha^n_i)_{0 < n < \omega} : i \in \mathcal{J} \rangle$ is indiscernible over $\{\mathbf{b}_n \wedge \beta_n : n < \omega\}$.

Step 3 : Conclusion

Claim : The sequences $(\mathbf{b}_n \wedge \beta_n)_{n < \omega}$ and $\langle (\mathbf{a}^n_i \wedge \alpha^n_i)_{i \in \mathcal{J}} : 0 < n < \omega \rangle$ are weakly linked. The sequences $(\mathbf{b}_n \wedge \beta_n)_{n < \omega}$ and \mathbf{U} are mutually indiscernible.

Proof. For the first statement, the proof is the same as in Section 5.2.4.

The second statement is similar. If for example we have $\phi(\mathbf{b}_n, \beta_n, \alpha^n_i)$, then the formula $\phi(\mathbf{b}_n, \beta_n, \alpha^n_j)$ must hold for all $j \in \mathcal{J}$, and therefore by total indiscernibility of \mathbf{U} and NIP, $\phi(\mathbf{b}_n, \beta_n, \alpha^n_j)$ must hold for every $(j, m) \in \mathcal{J} \times \omega$. □

Let $(\mathbf{c}_n, \gamma_n) = (\mathbf{a}^n_{1-\frac{1}{n}}, \alpha^n_{1-\frac{1}{n}})$, then :

1. The sequence $(\mathbf{c}_n \wedge \gamma_n)_{n < \omega}$ is indiscernible, with same EM-type as \mathbf{I} ;

2. The sequences $(\gamma_n)_{n < \omega}$ and $(b_n \hat{\gamma}_n)_{n < \omega}$ are mutually indiscernible;
3. The sequences $(c_n \hat{\gamma}_n)_{n < \omega}$ and $(b_n \hat{\gamma}_n)_{n < \omega}$ are weakly linked;
4. We have $\text{tp}(c_n \hat{\gamma}_n, b_m) = r$ if and only if $n = m$.

Consider the indiscernible sequence $(c_n \hat{\gamma}_n b_n \hat{\gamma}_n)_{n < \omega}$. We may increase it to an indiscernible sequence $(c_n \hat{\gamma}_n b_n \hat{\gamma}_n)_{n \in \mathcal{J}}$. Take some $n_0 \in \mathcal{J}$ and set $\mathcal{J} = \mathcal{J}_1 + \{n_0\} + \mathcal{J}_2$. Then by point 3 above, the sequence $\langle b_n \hat{\gamma}_n : n \in \mathcal{J}_1 + \mathcal{J}_2 \rangle$ is indiscernible over $c_{n_0} \hat{\gamma}_{n_0}$. Therefore point 4 and the definition of decomposition imply that β_{n_0} does not realize the limit type of $\langle \beta_n : n \in \mathcal{J}_1 \rangle$ over $\{b_n \hat{\gamma}_n : n \in \mathcal{J}_1 + \mathcal{J}_2\} \cup \{c_{n_0} \hat{\gamma}_{n_0}\}$. Adding parameters to the base, we may assume that it does not realize that limit type over $c_{n_0} \hat{\gamma}_{n_0}$.

We then iterate the construction, starting with the sequence $(c_n \hat{\gamma}_n b_n \hat{\gamma}_n)_{n \in \mathcal{J}}$. Assume that we can do this $|T|^+$ steps. We have at the end some base set of parameters A , an A -indiscernible sequence $\langle c_n \hat{\gamma}_n(\alpha_n^i : i < |T|^+) : n < \omega \rangle$ (we replaced the index set \mathcal{J} by ω for convenience) such that for each $i < |T|^+$, the sequence $(\alpha_n^i)_{n < \omega}$ is totally indiscernible over $A \cup \{\alpha_n^j : n < \omega, j \neq i\}$ but not indiscernible over $A \cup \{c_n, n < \omega\} \cup \{\alpha_n^j : n < \omega, j < i\}$. By Fodor's lemma, removing some sequences $(\alpha_n^i)_{n < \omega}$ and adding them to A , we may assume that for every i , $(\alpha_n^i)_{n < \omega}$ is not indiscernible over $A \cup \{d_n, n < \omega\}$. But this contradicts Proposition 5.1.17.

Therefore this construction must stop after less than $|T|^+$ stages. At the end, we obtain a decomposition of the sequence we started with. This proves Proposition 5.4.10.

Corollary 5.4.11. *Every dp-minimal theory is sharp.*

EXAMPLE 5.4.12 (Non-sharp theory). Let L_0 be the language $\{R_n(x, y) : n < \omega\}$ and construct an L_0 structure M_0 as follows : the universe of M_0 is \mathbb{Q} , the ordinary rational numbers, and for every $x, y \in M_0$, $M_0 \models R_n(x, y)$ if and only if $x < y \wedge |x - y| < n$ holds in \mathbb{Q} . Let $T_0 = \text{Th}(M_0)$. Non-realized 1-types over M_0 satisfying $R_n(x, a)$ for some $n < \omega$ and $a \in M$ are in natural bijection with cuts of $(\mathbb{Q}, <)$. In addition to these, there is just one non-realized type $p \in S_1(M_0)$ which satisfies $\neg R_n(x, a)$ for every $n < \omega$ and $a \in M$. This type p is generically stable (and \emptyset -invariant). One can check easily that T_0 is dp-minimal.

Now consider $L_1 = L_0 \cup \{<\}$ where $<$ is a new binary relation. We expand M_0 to an L_1 -structure M_1 by making $<$ into a generic order (i.e., every L_0 -infinite definable set of M_1 is dense co-dense with respect to $<$. See for example [59]). A 1-type over M_1 is determined by its reduct to L_0 plus its $<$ -cut. Let $T_1 = \text{Th}(M_1)$. Easily, T_1 eliminates imaginaries so there are no generically stable types (because the structure is linearly ordered). However T_1 is not distal : consider $I = (a_i)_{i \in \mathcal{J}}$ to be a dense $<$ -increasing sequence of points such that $\neg R_n(a_i, a_j)$ holds for every $n < \omega$ and $i, j \in \mathcal{J}$. Then this sequence is indiscernible and not distal. To see this, take two cuts c_1 and c_2 of I . Then there is a filling c_1 and b filling c_2 such that $R_1(a, b)$ holds. The generically stable type p in the reduct is detected by the non-distality of I .

We see however, that there is a natural ultra-imaginary stable sort : the quotient of M by the \vee -definable relation $E = \bigvee_{n < \omega} R_n$. And every point is in some sense s -dominated by its definable closure in that sort. It would be interesting to know if something like this is always true.

Théories dp-minimales ordonnées

§6.1 Introduction

One of the latest topics of interest in pure model theory is the study of dependent, or NIP, theories. The abstract general study was initiated by Shelah in [55], and pursued by him in [56], [53] and [58]. One of the questions he addresses is the definition of *super-dependent* as an analog of superstable for stable theories. Although, as he writes, he has not completely succeeded, the notion he defines of strong dependence seems promising. In [53] it is studied in detail and in particular, ranks are defined. Those so-called dp-ranks are used to prove existence of an indiscernible sub-sequence in any long enough sequence. Roughly speaking, a theory is strongly dependent if no type can fork infinitely many times, each forking being independent from the previous one. (Stated this way, it is naturally a definition of “strong NTP_2 ”). Also defined in that paper are notions of minimality, corresponding to the ranks being equal to 1 on 1-types. In [45], Onshuus and Usvyatsov extract from this material the notion of dp-minimality which seems to be the relevant one. A dp-minimal theory is a theory where there cannot be two independent witnesses of forking for a 1-type. It is shown in that paper that a stable theory is dp-minimal if and only if every 1-type has weight 1. In general, unstable, theories, one can link dp-minimality to *burden* as defined by H. Adler ([1]).

Dp-minimality on ordered structures can be viewed as a generalization of weak o-minimality. In that context, there are two main questions to address : what do definable sets in dimension 1 look like, (*i.e.* how far is the theory from being o-minimal), and what theorems about o-minimality go through. J. Goodrick has started to study those questions in [21], focussing on groups. He proves that definable functions are piecewise locally monotonous extending a similar result from weak-o-minimality.

In the first section of this paper, we recall the definitions and give equivalent formulations. In the second section, we make a few observations on general linearly ordered

inp-minimal theories showing in particular that, in dimension 1, forking is controlled by the ordering. The lack of a cell-decomposition theorem makes it unclear how to generalize results to higher dimensions.

In section 3, we study dp-minimal groups and show that they are abelian-by-finite-exponent. The linearly ordered ones are abelian. We prove also that an infinite definable set in a dp-minimal ordered divisible group has non-empty interior, solving a conjecture of A. Dolich.

Finally, in section 4, we give examples of dp-minimal theories. We prove that colored linear orders, orders of finite width and trees are dp-minimal.

I would like to thank John Goodrick and Alf Dolich for introducing me to some of the questions addressed in this paper. I also wish to thank Elisabeth Bouscaren and the referee for their thorough reading of the paper and for suggesting various improvements.

§6.2 Preliminaries on dp-minimality

Definition 6.2.1. (Shelah) An independence (or inp-) pattern of length κ is a sequence of pairs $(\phi^\alpha(x, y), k^\alpha)_{\alpha < \kappa}$ consisting of a formula and a positive integer such that there exists an array $\langle a_i^\alpha : \alpha < \kappa, i < \lambda \rangle$ for some $\lambda \geq \omega$ such that :

- Rows are k^α -inconsistent : for each $\alpha < \kappa$, the set $\{\phi^\alpha(x, a_i^\alpha) : i < \lambda\}$ is k^α -inconsistent,
- paths are consistent : for all $\eta \in \lambda^\kappa$, the set $\{\phi^\alpha(x, a_{\eta(\alpha)}^\alpha) : \alpha < \kappa\}$ is consistent.

Definition 6.2.2. – (Goodrick) A theory is inp-minimal if there is no inp-pattern of length two in a single free variable x .

- (Onshuus and Usvyatsov) A theory is dp-minimal if it is NIP and inp-minimal.

A theory is NTP_2 if there is no inp-pattern of size ω for which the formulas $\phi^\alpha(x, y)$ in the definition above are all equal to some $\phi(x, y)$. It is proven in [14] that a theory is NTP_2 if this holds for formulas $\phi(x, y)$ where x is a single variable. As a consequence, any inp-minimal theory is NTP_2 .

We now give equivalent definitions (all the ideas are from [53], we merely adapt the proofs there from the general NIP context to the dp-minimal one).

Definition 6.2.3. Two sequences $(a_i)_{i \in I}$ and $(b_j)_{j \in J}$ are *mutually indiscernible* if each one is indiscernible over the other.

Lemma 6.2.4. Consider the following statements :

1. T is inp-minimal.
2. For any two mutually indiscernible sequences $A = (a_i : i < \omega)$, $B = (b_j : j < \omega)$ and any point c , one of the sequences $(\text{tp}(a_i/c) : i < \omega)$, $(\text{tp}(b_i/c) : i < \omega)$ is constant.

3. Same as above, but change the conclusion to : one the sequences A or B stays indiscernible over c .
4. For any indiscernible sequence $A = (a_i : i \in I)$ indexed by a dense linear order I , and any point c , there is i_0 in the completion of I such that the two sequences $(tp(a_i/c) : i < i_0)$ and $(tp(a_i/c) : i > i_0)$ are constant.
5. Same as above, but change the conclusion to : the two sequences $(a_i : i < i_0)$ and $(a_i : i > i_0)$ are indiscernible over c .
6. T is dp-minimal.

Then for any theory T , (2), (3), (4), (5), (6) are equivalent and imply (1). If T is NIP, then they are all equivalent.

Proof. (2) \Rightarrow (1) : In the definition of independence pattern, one may assume that the rows are mutually indiscernible. This is enough.

(2) \Rightarrow (3) : Assume $A = \langle a_i : i < \omega \rangle$, $B = \langle b_i : i < \omega \rangle$ and c are a witness to $\neg(3)$. Then there are two tuples $(i_1 < \dots < i_n)$, $(j_1 < \dots < j_n)$ and a formula $\phi(x; y_1, \dots, y_n)$ such that $\models \phi(c; a_{i_1}, \dots, a_{i_n}) \wedge \neg \phi(c; a_{j_1}, \dots, a_{j_n})$. Take an $\alpha < \omega$ greater than all the i_k and the j_k . Then, exchanging the i_k and j_k if necessary, we may assume that $\models \phi(c; a_{i_1}, \dots, a_{i_n}) \wedge \neg \phi(c; a_{n.\alpha}, \dots, a_{n.\alpha+n-1})$. Define

$$A' = \langle (a_{i_1}, \dots, a_{i_n}) \rangle^\wedge \langle (a_{n.k}, \dots, a_{n.k+n-1}) : k \geq \alpha \rangle.$$

Construct the same way a sequence B' . Then A' , B' , c give a witness of $\neg(2)$.

(3) \Rightarrow (2) : Obvious.

(3) \Rightarrow (5) : Let $A = \langle a_i : i \in I \rangle$ be indiscernible and let c be a point. Then assuming (3) holds, for every i_0 in the completion of I , one of the two sequences $A_{<i_0} = \langle a_i : i < i_0 \rangle$ and $A_{>i_0} = \langle a_i : i > i_0 \rangle$ must be indiscernible over c . Take any such i_0 such that both sequences are infinite, and assume for example that $A_{>i_0}$ is indiscernible over c . Let $j_0 = \inf\{i \leq i_0 : A_{>i} \text{ is indiscernible over } c\}$. Then $A_{>j_0}$ is indiscernible over c . If there are no elements in I smaller than j_0 , we are done. Otherwise, if $A_{<j_0}$ is not indiscernible over c , then one can find $j_1 < j_0$ such that again $A_{<j_1}$ is not indiscernible over c . By definition of j_0 , $A_{>j_1}$ is not indiscernible over c either. This contradicts (3).

(5) \Rightarrow (4) : Obvious.

(4) \Rightarrow (2) : Assume $\neg(2)$. Then one can find a witness of it consisting of two indiscernible sequences $A = \langle a_i : i \in I \rangle$, $B = \langle b_i : i \in I \rangle$ indexed by a dense linear order I and a point c .

Now, we can find an i_0 in the completion of I such that for any $i_1 < i_0 < i_2$ in I , there are i, i' , $i_1 < i < i_0 < i' < i_2$ such that $tp(a_i/c) \neq tp(a_{i'}/c)$. Find a similar point j_0 for the sequence B . Renumbering the sequences if necessary, we may assume that $i_0 \neq j_0$. Then the indiscernible sequence of pairs $\langle (a_i, b_i) : i \in I \rangle$ gives a witness of $\neg(4)$.

(6) \Rightarrow (2) : Let A , B , c be a witness of $\neg(2)$. Assume for example that there is $\phi(x, y)$ such that $\models \phi(c, a_0) \wedge \neg \phi(c, a_1)$. Then set $A' = \langle (a_{2k}, a_{2k+1}) : k < \omega \rangle$ and $\phi'(x, y_1, y_2) = \phi(x, y_1) \wedge \neg \phi(x, y_2)$. Then by NIP, the set $\{\phi'(x, \bar{y}) : \bar{y} \in A'\}$ is k -inconsistent for some k . Doing the same construction with B we see that we get an independence pattern of length 2.

(5) \Rightarrow (6) : Statement (5) clearly implies NIP (because IP is always witnessed by a formula $\phi(x, y)$ with x a single variable). We have already seen that it implies inp-minimality. \square

Standard examples of dp-minimal theories include :

- O-minimal or weakly o-minimal theories (recall that a theory is weakly-o-minimal if every definable set in dimension 1 is a finite union of convex sets),
- C-minimal theories,
- $\text{Th}(\mathbf{Z}, +, \leq)$,
- The theory of the p -adics.

We refer the reader to [18] for more details and some proofs.

More examples are given in section 4 of this paper.

§6.3 Inp-minimal ordered structures

Little study has been made yet on general dp-minimal ordered structures. We believe however that there are results to be found already at that general level. In fact, we prove here a few lemmas that turn out to be useful for the study of groups.

We show that, in some sense, forking in dimension 1 is controlled by the order.

We consider $(M, <)$ an inp-minimal linearly ordered structure with no first nor last element. Let \mathfrak{C} be a monster model of $\text{Th}(M)$.

Lemma 6.3.1. *Let $X = X_{\bar{a}}$ be a definable subset of \mathfrak{C} , cofinal in \mathfrak{C} . Then X is non-forking (over \emptyset).*

Proof. If $X_{\bar{a}}$ divides over \emptyset , there exists an indiscernible sequence $(\bar{a}_i)_{i < \omega}$, $\bar{a}_0 = \bar{a}$, witnessing this. Every $X_{\bar{a}_i}$ is cofinal in \mathfrak{C} . Now pick by induction intervals I_k , $k < \omega$, with $I_k < I_{k+1}$ containing a point in each $X_{\bar{a}_i}$. We obtain an inp-pattern of length 2 by considering $x \in X_{\bar{a}_i}$ and $x \in I_k$.

If $X_{\bar{a}}$ forks over \emptyset , it implies a disjunction of formulas that divide, but one of these formulas must be cofinal : a contradiction. \square

A few variations are possible here. For example, we assumed that X was cofinal in the whole structure \mathfrak{C} , but the proof also works if X is cofinal in a \emptyset -definable set Y , or even contains an \emptyset -definable point in its closure. This leads to the following results.

For X a definable set, let $\text{Conv}(X)$ denote the convex hull of X . It is again a definable set.

Porism 6.3.2. *Let X be a definable set of \mathfrak{C} (in dimension 1). Assume $\text{Conv}(X)$ is A -definable. Then X is non-forking over A .*

Porism 6.3.3. *Let $M \prec N$ and let p be a complete 1-type over N . If the cut of p over N is of the form $+\infty$, $-\infty$, a^+ or a^- for $a \in M$, then p is non-forking over M .*

Proposition 6.3.5 generalizes this.

Lemma 6.3.4. *Let X be an A -definable subset of \mathfrak{C} . Assume that X divides over some model M , then :*

1. *We cannot find $(a_i)_{i < \omega}$ in M and points $(x_i)_{i < \omega}$ in $X(\mathfrak{C})$ such that $a_0 < x_0 < a_1 < x_1 < a_2 < \dots$.*
2. *The set X can be written as a finite disjoint union $X = \bigcup X_i$ where the X_i are definable over $M \cup A$, and each $\text{Conv}(X_i)$ contains no M -point.*

Proof. Easy ; (2) follows from (1). \square

Proposition 6.3.5. *Let $A \subset M$, with M , $|A|^+$ -saturated, and let $p \in S_1(M)$. The following are equivalent :*

1. *The type p forks over A ,*
2. *There exist $a, b \in M$ such that $p \vdash a < x < b$, and a and b have the same type over A ,*
3. *There exist $a, b \in M$ such that $p \vdash a < x < b$, and the interval $I_{a,b} = \{x : a < x < b\}$ divides over A .*

Proof. (3) \Rightarrow (1) is trivial.

For (2) \Rightarrow (3), it is enough to show that if $a \equiv_A b$, then $I_{a,b}$ divides over A . Let σ be an A -automorphism sending a to b . Then the tuple $(b = \sigma(a), \sigma(b))$ has the same type as (a, b) , and $a < b < \sigma(b)$. By iterating, we obtain a sequence $a_1 < a_2 < \dots$ such that (a_k, a_{k+1}) has the same type over A as (a, b) . Now the sets $I_{a_{2k}, a_{2k+1}}$ are pairwise disjoint and all have the same type over A . Therefore each of them divides over M .

We now prove (1) \Rightarrow (2)

Assume that (2) fails for p . Let $X_{\bar{a}}$ be an M -definable set such that $p \vdash X_{\bar{a}}$. Let $\bar{a}_0 = a, \bar{a}_1, \bar{a}_2, \dots$ be an A -indiscernible sequence. Note that the cut of p is invariant under all A -automorphisms. Therefore each of the $X_{\bar{a}_i}$ contains a type with the same cut over M as p . Now do a similar reasoning as in Lemma 6.3.1. \square

Corollary 6.3.6. *Forking equals dividing : for any $A \subset B$, any $p \in S(B)$, p forks over A if and only if p divides over A .*

Proof. By results of Chernikov and Kaplan ([15]), it is enough to prove that no type forks over its base. And it suffices to prove this for one-types (because of the general fact that if $\text{tp}(a/B)$ does not fork over A and $\text{tp}(b/Ba)$ does not fork over Aa , then $\text{tp}(a, b/B)$ does not fork over A).

Assume $p \in S_1(A)$ forks over A . Then by the previous proposition, p implies a finite disjunction of intervals $\bigcup_{i < n} (a_i, b_i)$ with $a_i \equiv_A b_i$. Assume n is minimal. Without loss, assume $a_0 < a_1 < \dots$. Now, as $a_0 \equiv_A b_0$ we can find points a'_i, b'_i , with $(a_i, b_i) \equiv_A (a'_i, b'_i)$ and $a'_0 = b_0$.

Then p proves $\bigcup_{i < n} (a'_i, b'_i)$. But the interval (a_0, b_0) is disjoint from that union, so p proves $\bigcup_{0 < i < n} (a_i, b_i)$, contradicting the minimality of n . \square

Note that this does not hold without the assumption that the structure is linearly ordered. In fact the standard example of the circle with a predicate $C(x, y, z)$ saying that y is between x and z (see for example [69], 2.2.4.) is dp-minimal.

Lemma 6.3.7. *Let E be a definable equivalence relation on M , we consider the imaginary sort $S = M/E$. Then there is on S a definable equivalence relation \sim with finite classes such that there is a definable linear order on S/\sim .*

Proof. Define a partial order on S by $a/E \prec b/E$ if $\inf(\{x : xEa\}) < \inf(\{x : xEb\})$. Let \sim be the equivalence relation on S defined by $x \sim y$ if $\neg(x \prec y \vee y \prec x)$. Then \prec defines a linear order on S/\sim . The proof that \sim has finite classes is another variation on the proof of 6.3.1. \square

From now until the end of this section, we also assume NIP.

Recall that in an NIP theory, if a type p splits over a model M , then it forks over M . In other words, if a, a' have the same type over M , then the formula $\phi(x, a) \Delta \phi(x, a')$ forks over M . (Note that the converse : “if p forks over M , then it splits over M ” is true in any theory.)

Lemma 6.3.8. (NIP). *Let $p \in S_1(\mathfrak{C})$ be a type inducing an M -definable cut, then p is definable over M .*

Proof. We know that p does not fork over M , so by NIP, p does not split over M . Let M_1 be an $|M|^+$ -saturated model containing M . Then the restriction of p to M_1 has a unique M -invariant extension. Therefore by NIP, it has a unique global extension that does not fork over M . This in turn implies by 6.3.5 that $p|_{M_1}$ has a unique global extension inducing the same cut as p , in particular it has a unique heir.

Therefore p is definable, and being M -invariant, p is definable over M . \square

The next lemma states that members of a uniformly definable family of sets define only finitely many “germs at $+\infty$ ”.

Lemma 6.3.9. (NIP). *Let $\phi(x, y)$ be a formula with parameters in some model M_0 , x a single variable. Then there are b_1, \dots, b_n such that for every b , there is $\alpha \in \mathfrak{C}$ and k such that the sets $\phi(x, b) \wedge x > \alpha$ and $\phi(x, b_k) \wedge x > \alpha$ are equal.*

Proof. Let E be the equivalence relation defined on tuples by bEb' iff $(\exists \alpha)(x > \alpha \rightarrow (\phi(x, b) \leftrightarrow \phi(x, b')))$. Let b, b' having the same type over M_0 . By NIP, the formula $\phi(x, b) \Delta \phi(x, b')$ forks over M_0 . By Lemma 6.3.1, this formula cannot be cofinal, so b and b' are E -equivalent. This proves that E has finitely many classes. \square

If the order is dense, then this analysis can be done also locally around a point a with the same proof :

Lemma 6.3.10. (NIP + dense order). *Let $\phi(x, y)$ be a formula with parameters in some model M_0 , x a single variable. Then there exists n such that : For any point a , there are b_1, \dots, b_n such that for all b , there is $\alpha < a < \beta$ and k such that the sets $\phi(x, b) \wedge \alpha < x < \beta$ and $\phi(x, b_k) \wedge \alpha < x < \beta$ are equal.*

§6.4 Dp-minimal groups

We study inp-minimal groups. Note that by an example of Simonetta, ([64]), not all such groups are abelian-by-finite. It is proven in [39] that C-minimal groups are abelian-by-torsion. We generalize the statement here to all inp-minimal theories.

Proposition 6.4.1. *Let G be an inp-minimal group. Then there is a definable normal abelian subgroup H such that G/H is of finite exponent.*

Proof. Let A, B be two definable subgroups of G . If $a \in A$ and $b \in B$, then there is $n > 0$ such that either $a^n \in B$ or $b^n \in A$. To see this, assume $a^n \notin B$ and $b^n \notin A$ for all $n > 0$. Then, for $n \neq m$, the cosets $a^n B$ and $a^m B$ are distinct, as are $A.b^m$ and $A.b^n$. Now we obtain an independence pattern of length two by considering the sequences of formulas $\phi_k(x) = "x \in a^k B"$ and $\psi_k(x) = "x \in A.b^k"$.

For $x \in G$, let $C(x)$ be the centralizer of x . By compactness, there is k such that for $x, y \in G$, for some $k' \leq k$, either $x^{k'} \in C(y)$ or $y^{k'} \in C(x)$. In particular, letting $n = k!$, x^n and y^n commute.

Let $H = C(C(G^n))$, the bicommutant of the n th powers of G . It is an abelian definable subgroup of G and for all $x \in G$, $x^n \in H$. Finally, if H contains all n powers then it is also the case of all conjugates of H , so replacing H by the intersection of its conjugates, we obtain what we want. \square

Now we work with ordered groups.

Note that in such a group, the convex hull of a subgroup is again a subgroup.

Lemma 6.4.2. *Let G be an inp-minimal ordered group. Let H be a definable subgroup of G and let C be the convex hull of H . Then H is of finite index in C .*

Proof. We may assume that H and C are \emptyset -definable. So without loss, assume $C = G$.

If H is not of finite index, there is a coset of H that forks over \emptyset . All cosets of H are cofinal in G . This contradicts Lemma 6.3.1. \square

Proposition 6.4.3. *Let G be an inp-minimal ordered group, then G is abelian.*

Proof. Note that if $a, b \in G$ are such that $a^n = b^n$, then $a = b$, for if for example $a < b$, then $a^n < a^{n-1}b < a^{n-2}b^2 < \dots < b^n$.

For $x \in G$, let $C(x)$ be the centralizer of x . We let also $D(x)$ be the convex hull of $C(x)$. By 6.4.2, $C(x)$ is of finite index in $D(x)$. Now take $x \in G$ and $y \in D(x)$. Then xy is in $D(x)$, so there is n such that $(xy)^n \in C(x)$. Therefore $(yx)^n = x^{-1}(xy)^n x = (xy)^n$. So $xy = yx$ and $y \in C(x)$. Thus $C(x) = D(x)$ is convex.

Now if $0 < x < y \in G$, then $C(y)$ is a convex subgroup containing y , so it contains x , and x and y commute. \square

This answers a question of Goodrick ([21] 1.1).

From now on, we will write groups additively.

Now, we assume NIP, so G is a dp-minimal ordered group. We denote by G^+ the set of positive elements of G .

Let $\phi(x)$ be a definable set (with parameters). For $\alpha \in G$, define $X_\alpha = \{g \in G^+ : (\forall x > \alpha)(\phi(x) \leftrightarrow \phi(x+g))\}$. Let H_α be equal to $X_\alpha \cup -X_\alpha \cup \{0\}$. Then H_α is a definable subgroup of G and if $\alpha < \beta$, H_α is contained in H_β . Finally, let H be the union of the H_α for $\alpha \in G$, it is the subgroup of *eventual periods* of $\phi(x)$.

Now apply Lemma 6.3.9 to the formula $\psi(x, y) = \phi(x-y)$. It gives n points b_1, \dots, b_n such that for all $b \in G$, there is k such that $b - b_k$ is in H . This implies that H has finite index in G .

If furthermore G is densely ordered, then we can do the same analysis locally. This yields a proof of a conjecture of A. Dolich : in a dp-minimal divisible ordered group, any infinite set has non empty interior. As a consequence, a dp-minimal divisible definably complete ordered group is o-minimal.

We will make use of two lemmas from [21] that we recall here for convenience.

Lemma 6.4.4 ([21], 3.3). *Let G be a densely ordered inp-minimal group, then any infinite definable set is dense in some non trivial interval.*

In the following lemma, \overline{G} stands for the completion of G . By a definable function f into \overline{G} , we mean a function of the form $a \mapsto \inf \phi(a; G)$ where $\phi(x; y)$ is a definable function. So one can view \overline{G} as a collection of imaginary sorts (in which case it naturally contains only *definable* cuts of G), or understand $f : G \rightarrow \overline{G}$ simply as a notation.

Lemma 6.4.5 (special case of [21], 3.19). *Let G be a densely ordered group, $f : G \rightarrow \overline{G}$ be a definable partial function such that $f(x) > 0$ for all x in the domain of f . Then for every interval I , there is a sub-interval $J \subseteq I$ and $\epsilon > 0$ such that for $x \in J \cap \text{dom}(f)$, $|f(x)| \geq \epsilon$.*

Theorem 6.4.6. *Let G be a divisible ordered dp-minimal group. Let X be an infinite definable set, then X has non-empty interior.*

Proof. As before, $I_{a,b}$ denotes the open interval $a < x < b$, and τ_b is the translation by $-b$.

Let $\phi(x)$ be a formula defining X .

By Lemma 6.4.4, there is an interval I such that X is dense in I . By Lemma 6.3.10 applied to $\psi(x; y) = \phi(y+x)$ at 0, there are $b_1, \dots, b_n \in M$ such that for all $b \in M$, there is $\alpha > 0$ and k such that $|x| < \alpha \rightarrow (\phi(b+x) \leftrightarrow \phi(b_k+x))$.

Taking a smaller I and X , if necessary, assume that for all $b \in I \cap X$, we may take $k = 1$.

Define $f : x \mapsto \sup\{y : I_{-y,y} \cap \tau_{b_1}X = I_{-y,y} \cap \tau_x X\}$, it is a function into \overline{M} , the completion of M . By Lemma 6.4.5, there is $J \subset I$ such that, for all $b \in J$, we have $|f(b)| \geq \epsilon$.

Fix $v < \frac{\epsilon}{2}$ and $b \in J$ such that $I_{b-2\epsilon, b+2\epsilon} \subseteq J$ (taking smaller ϵ if necessary). Set $L = I_{b-v, b+v}$ and $Z = L \cap X$. Assume for simplicity $b = 0$. Easily, if $g_1, g_2 \in Z$,

then $g_1 + g_2 \in Z \cup (G \setminus L)$ and $-g_1 \in Z$ (because the two points 0 and g_1 in Z have isomorphic neighborhoods of size ϵ). So Z is a group interval : it is the intersection with $I_{b-v, b+v}$ of some subgroup H of G . Now if $x, y \in L$ satisfy that there is $\alpha > 0$ such that $I_{-\alpha, \alpha} \cap \tau_x X = I_{-\alpha, \alpha} \cap \tau_y X$, then $x \equiv y$ modulo H . It follows that points of L lie in finitely many cosets modulo H . Assume Z is not convex, and take $g \in L \setminus Z$. Then for each $n \in \mathbf{N}$, the point g/n is in L and the points g/n define infinitely many different cosets ; a contradiction.

Therefore Z is convex and X contains a non trivial interval. \square

Corollary 6.4.7. *Let G be a dp-minimal ordered group. Assume G is divisible and definably complete, then G is o-minimal.*

Proof. Let X be a definable subset of G . By 6.4.6, the (topological) border Y of X is finite.

Let $a \in X$, then the largest convex set in X containing a is definable. By definable completeness, it is an interval and its end-points must lie in Y . As Y is finite, X is a finite union of (closed or open) intervals. \square

§6.5 Examples of dp-minimal theories

We give examples of dp-minimal theories, namely : linear orders, order of finite width and trees.

We first look at linear orders. We consider structures of the form (M, \leq, C_i, R_j) where \leq defines a linear order on M , the C_i are unary predicates (“colors”), the R_j are binary monotone relations (that is $x_1 \leq x R_j y \leq y_1$ implies $x_1 R_j y_1$).

The following is a (weak) generalization of Rubin’s theorem on linear orders (see [51], or [50]).

Proposition 6.5.1. *Let (M, \leq, C_i, R_j) be a colored linear order with monotone relations. Assume that all \emptyset -definable sets in dimension 1 are coded by a color and all monotone \emptyset -definable binary relations are represented by one of the R_j . Then the structure eliminates quantifiers.*

Proof. The result is obvious if M is finite, so we may assume (for convenience) that this is not the case.

We prove the theorem by back-and-forth. Assume that M is ω -saturated and take two tuples $\bar{x} = (x_1, \dots, x_n)$ and $\bar{y} = (y_1, \dots, y_n)$ from M having the same quantifier free type.

Take $x_0 \in M$; we look for a corresponding y_0 . Notice that \leq is itself a monotone relation, a finite boolean combinations of colors is again a color, a positive combination of monotone relations is again a monotone relation, and if $x R y$ is monotone $\phi(x, y) = \neg y R x$ is monotone. By compactness, it is enough to find a y_0 satisfying some finite part of the quantifier-free type of x_0 ; that is, we are given

- One color C such that $M \models C(x_0)$,
- For each k , monotone relations R_k and S_k such that $M \models x_0 R_k x_k \wedge x_k S_k x_0$.

Define $U_k(x) = \{t : t R_k x_k\}$ and $V_k(x) = \{t : x S_k t\}$. The $U_k(x)$ are initial segments of M and the $V_k(x)$ final segments. For each k, k' , either $U_k(x_k) \subseteq U_{k'}(x_{k'})$ or $U_{k'}(x_{k'}) \subseteq U_k(x_k)$. Assume for example $U_k(x_k) \subseteq U_{k'}(x_{k'})$, then this translates into a relation $\phi(x_k, x_{k'})$, where $\phi(x, y) = (\forall t)(t R_k x \rightarrow t R_{k'} y)$. Now $\phi(x, y)$ is a monotone relation itself. The assumptions on \bar{x} and \bar{y} therefore imply that also $U_k(y_k) \subseteq U_{k'}(y_{k'})$.

The same remarks hold for the final segments V_k .

Now, we may assume that $U_1(x_1)$ is minimal in the $U_k(x_k)$ and $V_1(x_1)$ is minimal in the $V_k(x_k)$. We only need to find a point y_0 satisfying $C(x)$ in the intersection $U_1(y_1) \cap V_1(y_1)$.

Let $\psi(x, y)$ be the relation $(\exists t)(C(t) \wedge t R_1 y \wedge x R_1 t)$. This is a monotone relation. As it holds for (x_0, x_1) , it must also hold for (y_0, y_1) , and we are done. \square

The following result was suggested, in the case of pure linear orders, by John Goodrick.

Proposition 6.5.2. *Let $\mathcal{M} = (M, \leq, C_i, R_j)$ be a linearly ordered infinite structure with colors and monotone relations. Then $\text{Th}(\mathcal{M})$ is dp-minimal.*

Proof. By the previous result, we may assume that $T = \text{Th}(\mathcal{M})$ eliminates quantifiers. Let $(x_i)_{i \in I}$, $(y_i)_{i \in I}$ be mutually indiscernible sequences of n -tuples, and let $\alpha \in M$ be a point. We want to show that one of the following holds :

- For all $i, i' \in I$, x_i and $x_{i'}$ have the same type over α , or
- for all $i, i' \in I$, y_i and $y_{i'}$ have the same type over α .

Assume that I is dense without end points.

By quantifier elimination, we may assume that $n = 1$, that is the x_i and y_i are points of M . Without loss, the (x_i) and (y_i) form increasing sequences. Assume there exists $i < j \in I$ and R a monotone definable relation such that $M \models \neg \alpha R x_i \wedge \alpha R x_j$. By monotonicity of R , there is a point i_R of the completion of I such that $i < i_R \rightarrow \neg \alpha R x_i$ and $i > i_R \rightarrow \alpha R x_i$.

Assume there is also a monotone relation S and an i_S such that $i < i_S \rightarrow \neg \alpha S y_i$ and $i > i_S \rightarrow \alpha S y_i$.

For points x, y define $I(x, y)$ as the set of $t \in M$ such that $M \models \neg t R x \wedge t R y$. This is an interval of M . Furthermore, if $i_1 < i_2 < i_3 < i_4$ are in I , then the intervals $I(x_{i_1}, x_{i_2})$ and $I(x_{i_3}, x_{i_4})$ are disjoint. Define $J(x, y)$ the same way using S instead of R .

Take $i_0 < i_R < i_1 < i_2 < \dots$ and $j_0 < i_S < j_1 < j_2 < \dots$. For $k < \omega$, define $I_k = I(x_{i_{2k}}, x_{i_{2k+1}})$ and $J_k = J(y_{j_{2k}}, y_{j_{2k+1}})$. The two sequences (I_k) and (J_k) are mutually indiscernible sequences of disjoint intervals. Furthermore, we have $\alpha \in I_0 \cap J_0$. By mutual indiscernibility, $I_i \cap J_j \neq \emptyset$ for all indices i and j , which is impossible.

We treated the case when α was to the left of the increasing relations R and S . The other cases are similar. \square

An ordered set (M, \leq) is of *finite width*, if there is n such that M has no antichain of size n .

Corollary 6.5.3. *Let $\mathcal{M} = (M, \leq)$ be an infinite ordered set of finite width, then $\text{Th}(\mathcal{M})$ is dp-minimal.*

Proof. We can define such a structure in a linear order with monotone relations : see [52]. More precisely, there exists a structure $P = (P, \prec, R_j)$ in which \prec is a linear order and the R_j are monotone relations, and there is a definable relation $O(x, y)$ such that the structure (P, O) is isomorphic to (M, \leq) .

The result therefore follows from the previous one. \square

We now move to trees. A tree is a structure (T, \leq) such that \leq defines a partial order on T , and for all $x \in T$, the set of points smaller than x is linearly ordered by \leq . We will also assume that given $x, y \in T$, the set of points smaller than x and y has a maximal element $x \wedge y$ (and set $x \wedge x = x$). This is not actually a restriction, since we could always work in an imaginary sort to ensure this.

Given $a, b \in T$, we define the open ball $B(a; b)$ of center a containing b as the set $\{x \in T : x \wedge b > a\}$, and the closed ball of center a as $\{x \in T : x \geq a\}$.

Notice that two balls are either disjoint or one is included in the other.

Lemma 6.5.4. *Let (T, \leq) be a tree, $a \in T$, and let D denote the closed ball of center a . Let $\bar{x} = (x^1, \dots, x^n) \in (T \setminus D)^n$ and $\bar{y} = (y^1, \dots, y^m) \in D^m$. Then $\text{tp}(\bar{x}/a) \cup \text{tp}(\bar{y}/a) \vdash \text{tp}(\bar{x} \cup \bar{y}/a)$.*

Proof. A straightforward back-and-forth, noticing that $\text{tp}(\bar{x}/a) \cup \text{tp}(\bar{y}/a) \vdash \text{tp}_{\text{qf}}(\bar{x} \cup \bar{y}/a)$ (quantifier-free type). \square

We now work in the language $\{\leq, \wedge\}$, so a sub-structure is a subset closed under \wedge .

Proposition 6.5.5. *Let $A = (a_0, \dots, a_n)$, $B = (b_0, \dots, b_n)$ be two sub-structures from T . Assume :*

1. *A and B are isomorphic as sub-structures,*
2. *for all i, j such that $a_i \geq a_j$, $\text{tp}(a_i, a_j) = \text{tp}(b_i, b_j)$.*

Then $\text{tp}(A) = \text{tp}(B)$.

Proof. We do a back-and-forth. Assume T is ω -saturated and A, B satisfy the hypothesis. We want to add a point a to A . We may assume that $A \cup \{a\}$ forms a sub-structure (otherwise, if some $a_i \wedge a$ is not in $A \cup \{a\}$, add first this element).

We consider different cases :

1. The point a is below all points of A . Without loss a_0 is the minimal element of A (which exists because A is closed under \wedge). Then find a b such that $\text{tp}(a_0, a) = \text{tp}(b_0, b)$. For any index i , we have : $\text{tp}(a_i, a_0) = \text{tp}(b_i, b_0)$ and $\text{tp}(a, a_0) = \text{tp}(b, b_0)$. By Lemma 6.5.4, $\text{tp}(a_i, a) = \text{tp}(b_i, b)$.

2. The point a is greater than some point in A , say a_1 , and the open ball $a := B(a_1; a)$ contains no point of A .

Let \mathcal{A} be the set of all open balls $B(a_1; a_i)$ for $a_i > a_1$. Let n be the number of balls in \mathcal{A} that have the same type p as a . Then $\text{tp}(a_1)$ proves that there are at least

$n + 1$ open balls of type p of center a_1 . Therefore, $\text{tp}(b_1)$ proves the same thing. We can therefore find an open ball b of center b_1 of type p that contains no point from B . That ball contains a point b such that $\text{tp}(b_1, b) = \text{tp}(a_1, a)$. Now, if a_i is smaller than a_1 , we have $\text{tp}(a_i, a_1) = \text{tp}(b_i, b_1)$ and $\text{tp}(a_1, a) = \text{tp}(b_1, b)$, therefore by Lemma 6.5.4, $\text{tp}(a, a_i) = \text{tp}(b, b_i)$.

The fact that we have taken b in a new open ball of center b_1 ensures that $B \cup \{b\}$ is again a sub-structure and that the two structures $A \cup \{a\}$ and $B \cup \{b\}$ are isomorphic.

3. The point a is between two points of A , say a_0 and a_1 ($a_0 < a_1$), and there are no points of A between a_0 and a_1 .

Find a point b such that $\text{tp}(a_0, a_1, a) = \text{tp}(b_0, b_1, b)$. Then if i is such that $a_i > a$, we have $a_i \geq a_1$ and again by Lemma 6.5.4, $\text{tp}(a_i, a) = \text{tp}(b_i, b)$. And same if $a_i < a$. \square

Corollary 6.5.6. *Let $A \subset T$ be any subset. Then $\bigcup_{(a,b,c) \in A^3} \text{tp}(a, b, c) \vdash \text{tp}(A)$.*

Proof. Let A_0 be the substructure generated by A . By the previous theorem the following set of formulas implies the type of A_0 :

- the quantifier-free type of A_0 ,
- the set of 2-types $\text{tp}(a, b)$ for $(a, b) \in A_0^2$, $a \leq b$.

We need to show that those formulas are implied by the set of 3-types of elements of A . We may assume A is finite.

First, the knowledge of all the 3-types is enough to construct the structure A_0 . To see this, start for example with a point $a \in A$ maximal. Knowing the 3-types, one knows in what order the $b \wedge a, b \in A$ are placed. Doing this for all such a , enables one to reconstruct the tree A_0 .

Now take $m_1 = a \wedge b$, $m_2 = c \wedge d$ for $a, b, c, d \in A$ such that $m_1 \leq m_2$. The points m_1 and m_2 are both definable using only 3 of the points a, b, c, d , say a, b, c . Then $\text{tp}(a, b, c) \vdash \text{tp}(m_1, m_2)$. \square

The previous results are also true, with the same proofs, for colored trees.

It is proven in [46] that theories of trees are NIP. We give a more precise result.

Proposition 6.5.7. *Let $\mathcal{T} = (T, \leq, C_i)$ be a colored tree. Then $\text{Th}(\mathcal{T})$ is dp-minimal.*

Proof. We will use criterion (3) of 6.2.4 : if $(a_i)_{i \in I}$ and $(b_j)_{j \in J}$ are mutually indiscernible sequences and $\alpha \in T$ is a point, then one of the sequences (a_i) and (b_j) is indiscernible over α .

We will always assume that the index sets $(I$ and $J)$ are dense linear orders without end points.

1) We start by showing the result assuming the a_i and b_j are points (not tuples).

We classify the indiscernible sequence (a_i) in 4 classes depending on its quantifier-free type.

I The sequence (a_i) is monotonous (increasing or decreasing).

II The α_i are pairwise incomparable and $\alpha_i \wedge \alpha_j$ is constant equal to some point β .

If (α_i) is in none of those two cases, consider indices $i < j < k$. Note that it is not possible that $\alpha_i \wedge \alpha_j < \alpha_i \wedge \alpha_k$, so there are two cases left to consider :

III $\alpha_i \wedge \alpha_k = \alpha_i \wedge \alpha_j$. Then let $\alpha'_i = \alpha_i \wedge \alpha_j$ (this does not depend on j , $j > i$). The α'_i form an increasing indiscernible sequence.

IV $\alpha_i \wedge \alpha_k < \alpha_i \wedge \alpha_j$. Then $\alpha'_j = \alpha_i \wedge \alpha_j$ is independent of the choice of i ($i < j$) and (α'_j) is a decreasing indiscernible sequence.

Assume (α_i) lands in case **I**. Consider the set $\{x : x < \alpha\}$. If that set contains a non-trivial subset of the sequence (α_i) , we say that α *cuts* the sequence. If this is not the case, then the sequence (α_i) stays indiscernible over α . To see this, assume for example that (α_i) is increasing and that α is greater than all the α_i . Take two sets of indices $i_1 < \dots < i_n$ and $j_1 < \dots < j_n$ and a $k \in I$ greater than all those indices. Then $\text{tp}(\alpha_{i_1}, \dots, \alpha_{i_n}/\alpha_k) = \text{tp}(\alpha_{j_1}, \dots, \alpha_{j_n}/\alpha_k)$. Therefore by Lemma 6.5.4, $\text{tp}(\alpha_{i_1}, \dots, \alpha_{i_n}/\alpha) = \text{tp}(\alpha_{j_1}, \dots, \alpha_{j_n}/\alpha)$.

In case **II**, note that if (α_i) is not α -indiscernible, then there is $i \in I$ such that α lies in the open ball $B(\beta; \alpha_i)$ (we will also say that α *cuts* the sequence (α_i)). This follows easily from Proposition 6.5.5.

In the last two cases, if (α_i) is α -indiscernible, then it is also the case for (α'_i) . Conversely, if (α'_i) is α -indiscernible, then α does not cut the sequence (α'_i) . From 6.5.5, it follows easily that (α_i) is also α -indiscernible. We can therefore replace the sequence (α_i) by (α'_i) which belongs to case **I**.

Going back to the initial data, we may assume that (α_i) and (b_j) are in case **I** or **II**. It is then straightforward to check that α cannot cut both sequences. For example, assume (α_i) is increasing and (b_j) is in case **II**. Then define β as $b_i \wedge b_j$ (any i, j). If α cuts (b_j) , then $\alpha > \beta$. But (α_i) is β -indiscernible. So β does not cut (α_i) . The only possibility for α to cut (α_i) is that β is smaller than all the α_i and the α_i lie in the same open ball of center β as α . But then the α_i lie in the same open ball of center β as one of the b_j . This contradicts mutual indiscernability.

2) Reduction to the previous case. We show that if $(\alpha_i)_{i \in I}$ is an indiscernible sequence of n -tuples and $\alpha \in T$ such that (α_i) is not α -indiscernible, then there is an indiscernible sequence $(d_i)_{i \in I}$ of points of T in $\text{dcl}((\alpha_i))$ such that (d_i) is not α -indiscernible.

First, by 6.5.6, we may assume that $n = 2$. Write $\alpha_i = (b_i, c_i)$ and define $m_i = b_i \wedge c_i$. We again study different cases :

1. The m_i are all equal to some m .

As (α_i) is not α -indiscernible, necessarily, $\alpha > m$ and the ball $B(m; \alpha)$ contains one b_i (resp. c_i). Then take $d_i = b_i$ (resp. $d_i = c_i$) for all i .

2. The m_i are linearly ordered by $<$ and no b_i nor c_i is greater than all the m_i . Then the balls $B(m_i; b_i)$ and $B(m_i; c_i)$ contain no other point from $(b_i, c_i, m_i)_{i \in I}$. Then, α must cut the sequence (m_i) and one can take $d_i = m_i$ for all i .

3. The m_i are linearly ordered by $<$ and, say, each b_i is greater than all the m_i .

Then each ball $B(\mathfrak{m}_i; \mathfrak{a}_i)$ contains no other point from $(\mathfrak{b}_i, \mathfrak{c}_i, \mathfrak{m}_i)_{i \in I}$. If α cuts the sequence \mathfrak{m}_i , then again one can take $\mathfrak{d}_i = \mathfrak{m}_i$. Otherwise, take a point γ larger than all the \mathfrak{m}_i but smaller than all the \mathfrak{d}_i . Applying 6.5.4 with \mathfrak{a} there replaced by γ , we see that (\mathfrak{b}_i) cannot be α -indiscernible. Then take $\mathfrak{d}_i = \mathfrak{b}_i$ for all i .

4. The \mathfrak{m}_i are pairwise incomparable.

The sequence (\mathfrak{m}_i) lies in case **II**, **III** or **IV**. The open balls $B(\mathfrak{m}_i; \mathfrak{b}_i)$ and $B(\mathfrak{m}_i; \mathfrak{c}_i)$ cannot contain any other point from $(\mathfrak{b}_i, \mathfrak{c}_i, \mathfrak{m}_i)_{i \in I}$. Considering the different cases, one sees easily that taking $\mathfrak{d}_i = \mathfrak{m}_i$ will work.

This finishes the proof.

□

Bibliographie

- [1] Hans Adler. Strong theories, burden, and weight. *in preparation*.
- [2] Hans Adler. An introduction to theories without the independence property. *Archive of Mathematical Logic*, 2008. to appear.
- [3] Yerzhan Baisalov and Bruno Poizat. Paires de structures o-minimales. *Journal of Symbolic Logic*, 63 :570–578, 1998.
- [4] Bektur Baizhanov and John Baldwin. Local homogeneity. *Journal of Symbolic Logic*, 69(4) :1243–1260, 12 2004.
- [5] John Baldwin and Michael Benedikt. Stability theory, permutations of indiscernibles, and embedded finite models. *Transactions of the American Mathematical Society*, 352(11) :4937–4969, 11 2000.
- [6] Oleg Belegradek, Ya’acov Peterzil, and Frank Wagner. Quasi-o-minimal structures. *J. Symbolic Logic*, 65(3) :1115–1132, 2000.
- [7] Oleg Belegradek and Boris Zilber. The model theory of the field of reals with a subgroup of the unit circle. *J. Lond. Math. Soc. (2)*, 78(3) :563–579, 2008.
- [8] Alexander Berenstein. Lovely pairs and dense pairs of o-minimal structures. Preprint.
- [9] Alexander Berenstein, Alfred Dolich, and Alf Onshuus. The independence property in generalized dense pairs of structures. *J. Symbolic Logic*, 76(2) :391–404, 2011.
- [10] Anselm Blumer, A. Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *J. ACM*, 36 :929–965, October 1989.
- [11] Gareth Boxall. Nip for some pair-like theories. *Archive for Mathematical Logic*, 50 :353–359, 2011. 10.1007/s00153-010-0218-3.
- [12] Enrique Casanovas and Martin Ziegler. Stable theories with a new preicate. *Journal of symbolic logic*, 66(3) :1127–1140, 09 2001.

- [13] Zoé Chatzidakis and Ehud Hrushovski. Model theory of difference fields. *Trans. Amer. Math. Soc.*, 351(8) :2997–3071, 1999.
- [14] Artem Chenikov. Theories without the tree property of the second kind. in preparation.
- [15] Artem Chernikov and Itay Kaplan. Forking and dividing in NTP_2 theories. accepted to the JSL, 2009.
- [16] Artem Chernikov and Pierre Simon. Externally definable sets and dependent pairs. Submitted, Modnet preprint 265.
- [17] Françoise Delon. Définissabilité avec paramètres extérieurs dans \mathbf{Q}_p et \mathbf{R} . *Proc. Amer. Math. Soc.*, 106(1) :193–198, 1989.
- [18] Alfred Dolich, David Lippel, and John Goodrick. Dp-minimal theories : basic facts and examples. Submitted.
- [19] Antonio J. Engler and Alexander Prestel. *Valued fields*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2005.
- [20] Antongiulio Fornasiero. Tame structures and open cores. Submitted, 2010.
- [21] John Goodrick. A monotonicity theorem for dp-minimal densely ordered groups. *J. Symb. Log.*, 75(1) :221–238, 2010.
- [22] Vincent Guingona. Dependence and isolated extensions. Modnet preprint 212, 2009.
- [23] Vincent Guingona. On uniform definability of types over finite sets. Modnet preprint 250, 2010.
- [24] Ayhan Günaydin and Philipp Hieronymi. The real field with the rational points of an elliptic curve. Modnet preprint 187.
- [25] Ayhan Günaydin and Philipp Hieronymi. Dependent pairs. Modnet preprint 146, 2008.
- [26] Bradd Hart, Byunghan Kim, and Anand Pillay. Coordinatisation and canonical bases in simple theories. *J. Symbolic Logic*, 65(1) :293–309, 2000.
- [27] Deirdre Haskell, Ehud Hrushovski, and Dugald Macpherson. *Stable domination and independence in algebraically closed valued fields*, volume 30 of *Lecture Notes in Logic*. Association for Symbolic Logic, Chicago, IL, 2008.
- [28] Wilfrid Hodges. *Model Theory*, volume 42 of *Encyclopedia of mathematics and its applications*. Cambridge University Press, Great Britain, 1993.
- [29] Ehud Hrushovsk, Anand Pillay, and Pierre Simon. Generically stable and smooth measures in NIP theories. Submitted.
- [30] Ehud Hrushovski. Pseudo-finite fields and related structures. In *Model theory and applications*, volume 11 of *Quad. Mat.*, pages 151–212. Aracne, Rome, 2002.
- [31] Ehud Hrushovski and Francois Loeser. Non-archimedean tame topology and stably dominated types. 2010.
- [32] Ehud Hrushovski, Ya’acov Peterzil, and Anand Pillay. Groups, measures, and the NIP. *J. Amer. Math. Soc.*, 21(2) :563–596, 2008.

- [33] Ehud Hrushovski and Anand Pillay. On nip and invariant measures. *Journal of the European Mathematical Society*, 13 :1005–1061, 2011.
- [34] H. R. Johnson and M. C. Laskowski. Compression schemes, stable definable families, and o-minimal structures. *Discrete Comput. Geom.*, 43(4) :914–926, 2010.
- [35] Marek Karpinski and Angus Macintyre. Approximating volumes and integrals in o-minimal and p-minimal theories. In *Connections between model theory and algebraic and analytic geometry*, volume 6 of *Quad. Mat.*, pages 149–177. Dept. Math., Seconda Univ. Napoli, Caserta, 2000.
- [36] H. Jerome Keisler. Choosing elements in a saturated model. In *Classification theory (Chicago, IL, 1985)*, volume 1292 of *Lecture Notes in Math.*, pages 165–181. Springer, Berlin, 1987.
- [37] H. Jerome Keisler. Measures and forking. *Ann. Pure Appl. Logic*, 34(2) :119–169, 1987.
- [38] Byunghan Kim and Anand Pillay. Simple theories. *Ann. Pure Appl. Logic*, 88(2-3) :149–164, 1997. Joint AILA-KGS Model Theory Meeting (Florence, 1995).
- [39] Dugald Macpherson and Charles Steinhorn. On variants of o-minimality. *Ann. Pure Appl. Logic*, 79(2) :165–209, 1996.
- [40] David Marker and Charles I. Steinhorn. Definable types in o-minimal theories. *J. Symbolic Logic*, 59(1) :185–198, 1994.
- [41] Jiri Matousek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.
- [42] Michael Morley. Categoricity in power. *Transaction of the American Mathematical Society*, 114 :514–538, 1965.
- [43] A. Onshuus and A. Pillay. Definable groups and compact p-adic Lie groups. *J. Lond. Math. Soc. (2)*, 78(1) :233–247, 2008.
- [44] Alf Onshuus and Yaacov Peterzil. A note on stable sets, groups, and theories with nip. *Mathematical Logic Quarterly*, 53, 2007.
- [45] Alf Onshuus and Alexander Usvyatsov. On dp-minimality, strong dependence and weight. *submitted*.
- [46] Michel Parigot. Théories d’arbres. *J. Symbolic Logic*, 47(4) :841–853 (1983), 1982.
- [47] A. Pillay. *Geometric stability theory*. Oxford logic guides. Clarendon Press, 1996.
- [48] Anand Pillay. On externally definable sets and a theorem of shelah. *Algebra, logic, set theory, Stud. Log. (Lond.)*, 4 :175–181, 2007.
- [49] Bruno Poizat. Paires de structures stables. *J. Symbolic Logic*, 48(2) :239–249, 1983.
- [50] Bruno Poizat. *Cours de théorie des modèles*. Nur al-Mantiq wal-Mari’fah, 1985.
- [51] Matatyahu Rubin. Theories of linear order. *Isr. J. Math.*, 17 :392–443, 1974.
- [52] James H. Schmerl. Partially ordered sets and the independence property. *J. Symb. Log.*, 54(2) :396–401, 1989.

- [53] Saharon Shelah. Strongly dependent theories. Submitted.
- [54] Saharon Shelah. *Classification theory and the number of nonisomorphic models*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, second edition, 1990.
- [55] Saharon Shelah. Classification theory for elementary classes with the dependence property—a modest beginning. *Sci. Math. Jpn.*, 59(2) :265–316, 2004. Special issue on set theory and algebraic model theory.
- [56] Saharon Shelah. Dependent first order theories, continued. *Israel Journal of Mathematics*, 173(1) :1–60, 2009.
- [57] Saharon Shelah. A dependent dream and recounting types. *preprint*, 2010. 950 in Shelah archive.
- [58] Saharon Shelah. Dependent theories and the generic pair conjecture. *Communications in Contemporary Mathematics*, submitted, 2010. math.LO/0702292.
- [59] Saharon Shelah and Pierre Simon. Adding linear orders. Submitted, Modnet preprint 290, 2011.
- [60] Pierre Simon. Théories nip. Master’s thesis, École normale supérieure, 2009. M. Sc. Thesis.
- [61] Pierre Simon. Distal and non-distal theories. 2011.
- [62] Pierre Simon. Finding generically stable measures. to appear in the Journal of Symbolic Logic, 2011.
- [63] Pierre Simon. On dp-minimal ordered structures. *Journal of Symbolic Logic*, 76(4) :448–460, 2011.
- [64] Patrick Simonetta. An example of a c-minimal group which is not abelian-by-finite. *Proc. Am. Math. Soc.*, 131(12) :3913–3917, 2003.
- [65] Alexander Usvyatsov. On generically stable types in dependent theories. *Journal of Symbolic Logic*. accepted.
- [66] Lou van den Dries. Dense pairs of o-minimal structures. *Fund. Math.*, 157 :61–78, 1998.
- [67] Lou van den Dries and Adam H. Lewenberg. T-convexity and tame extensions. *Journal of Symbolic Logic*, 155(3) :807–836, 1995.
- [68] V. N. Vapnik and A. Ya. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16(2) :264–280, 1971.
- [69] Frank O. Wagner. *Simple theories*, volume 503 of *Mathematics and its Applications*. Kluwer Academic Publishers, Dordrecht, 2000.
- [70] I. Ben Yaacov and H. J. Keisler. Randomization of models as metric structures. *Confluentes Mathematici* 1, (2) :197–223, 2009.
- [71] Itai Ben Yaacov. Continuous and random vapnik-chervonenkis classes. *Israel Journal of Mathematics*, 173 :309–333, 2009.

- [72] Itai Ben Yaacov. Transfer of properties between measures and random types. in preparation, 2009.