Università di Pisa
Dipartimento di Matematica
Corso di Laurea Magistrale in Matematica

Definable Groups, NIP Theories, and the Ellis Group Conjecture

Candidato	Relatore	Controrelatore
Rosario Mennuni	**Prof. Alessandro Berarducci**	**Prof. Mauro Di Nasso**

Anno Accademico 2014/2015

a Federica

Tra il dire e il fare c'è di mezzo "e il".
Elio e le Storie Tese

Contents

Introduction

Abstract

A *definable group* G is a group which is definable in a first-order structure. Despite the name, it is not a single group, but a family of groups given by interpreting the defining formulas in *elementary extensions* of the structure defining the group. For instance, algebraic groups are definable in the complex field using first-order formulas. These include matrix groups and abelian varieties such as elliptic curves. Among groups which are definable with first-order formulas in the real field there are $\mathrm{GL}(n, \mathbb{R})$, $\mathrm{SO}(n, \mathbb{R})$, and other Lie groups.

The two families of examples above are, in a sense, orthogonal. The field $\mathbb{C}$ falls into the class of *stable* structures, which are, in a nutshell, the ones that do not define an order relation on an infinite set. Stable theories have been a central and fruitful topic in the model theory of the past decades (e.g. [Pil96, Bal88]), and there is a huge literature on stable groups (for instance [Poi01, BN94]). Unfortunately, since stability is destroyed by the presence of a total infinite order, the field structure on $\mathbb{R}$ lives outside this realm, and more generally *o-minimal* structures, another important class in which it is possible to provide a framework for *tame geometry* (see [vdD98]), are not stable. Model theorists have therefore tried to generalize methods from stability theory to broader contexts. One robust, simultaneous generalization of both stability and o-minimality is found in the class of *dependent*, or NIP theories. NIP structures can be roughly described as the ones that do not code a membership relation on an infinite set; this viewpoint is intimately connected to VC-dimension, a fundamental tool of statistical learning theory. This thesis explores a problem, which we are now going to outline, concerning the relation between two groups that can be attached to any group definable in a NIP structure.

In NIP theories, to every definable group is associated a concrete compact Hausdorff topological group called G/G^{00}. As an example, it can be proven that if G is a definably compact group definable over $\emptyset$ in a real closed field, for instance $\mathrm{SO}(3, M)$ for $M \succ \mathbb{R}$ an hyperreal field, then G/G^{00} is exactly $G(\mathbb{R})$, and the projection to G/G^{00} behaves like a "standard part" map. If G

is not compact then this may not be true, as in the case of $\mathrm{SL}(n, M)$ where G/G^{00} is trivial. In general (see [BOPP05]), for a group which is definable in an o-minimal structure, G/G^{00} is a real Lie group. As a stable example, if G is the additive group in the structure of the integers with sum (but without product), then G/G^{00} is isomorphic to $\hat{\mathbb{Z}} = \varprojlim \mathbb{Z}/n\mathbb{Z}$. All these isomorphisms preserve the topology, i.e. are isomorphisms of topological groups. This canonical quotient is the first protagonist of the problem studied in the thesis. In order to introduce the second one, some preliminary explanations are needed.

An important concept in the study of stable groups is the one of a *generic type*. Trying to find a well-behaved analogue in the unstable context, Newelski noticed that a certain notion, namely that of a *weak generic type*, is well understood when bringing topological dynamics into the picture*. In topological dynamics one is often interested in G-flows, actions of a group G on compact Hausdorff spaces by homeomorphisms; soon one turns the attention to the ones that have a dense orbit (G-ambits) and to the ones in which all orbits are dense (minimal flows). A very special G-flow is the universal G-ambit βG of ultrafilters on G: every G-ambit can be seen as a quotient of βG, and its minimal subflows enjoy a similar universal property. A "tame" counterpart of βG is the space $S_G(M)$ of types over a model M concentrating on G, i.e. the ultrafilters on definable subsets of $G(M)$, and one could develop a theory of *tame topological dynamics* ([GPP14, Pil13]) and hope for $S_G(M)$ to be universal with respect to *definable* $G(M)$-flows. Now, one important tool in the study of a G-flow X is its *enveloping semigroup* $E(X)$; it turns out that $\beta G \cong E(\beta G)$ and this equips the former with a semigroup structure. Once some technical obstacles are overcome, this construction can be carried out for $S_G(M)$ too, or at least for a certain bigger type space called $S_G^{\mathrm{ext}}(M)$.

Applying the theory of enveloping semigroups to $E(\beta G) \cong \beta G$ produces a certain family of sub-semigroups that are indeed groups, and furthermore all in the same isomorphism class: this is the *ideal group*, or *Ellis group* associated to the flow. Modulo the complications mentioned above, an Ellis group can also be associated to $S_G(M)$. Even if this may depend on M, a comparison with G/G^{00} can be made, and indeed the latter is always a quotient of the former, the projection π being the restriction of a certain natural map $S_G(M) \to G/G^{00}$. Since in stable groups a similar situation arises replacing the Ellis group with the subspace of generic types of $S_G(M)$, and in that case the relevant map is injective, the next question is: is this π an isomorphism?

*Briefly, in the dynamical context "generic" becomes "syndetic", and "weak generic" corresponds to "piecewise syndetic".

Even in tame context, this need not be the case: it was shown in [GPP15] that the Ellis group of $\mathrm{SL}(2,\mathbb{R})$ is the group with two elements, but its G/G^{00} is trivial. A property that is *not* satisfied by $\mathrm{SL}(2,\mathbb{R})$ is amenability: there is no finitely additive, left-translation-invariant probability measure defined on $\mathscr{P}(\mathrm{SL}(2,\mathbb{R}))$. Another group lacking amenability is $\mathrm{SO}(3,\mathbb{R})$; this is essentially the Banach-Tarski paradox. The reasons behind the non-amenability of these two groups are, however, different. If one searches for a left-translation-invariant (finitely additive) measure defined not on the whole power-set, but only on the Boolean algebra of *definable* subsets of $\mathrm{SO}(3,\mathbb{R})$, then such a measure *does* exist, and we say that $\mathrm{SO}(3,\mathbb{R})$ is *definably amenable*. A similar thing happens with free groups on at least two generators. This is due to the fact that the non-measurable sets arising from the Banach-Tarski paradox are very complicated, and certainly not definable in the first-order structure of $\mathbb{R}$, and so this kind of obstructions to amenability disappear when we only want a measure on an algebra of "simple" sets. On the contrary, $\mathrm{SL}(2,\mathbb{R})$ is not even definably amenable, thus being more inherently pathological under this point of view. In [CPS14] Pillay then proposed the *Ellis Group Conjecture*. Several special cases were proven in the same paper and, thereafter, the conjecture was proven true in [CSed] by Chernikov and Simon, hence we state it as a Theorem.

Theorem ([CSed, Theorem 5.6])**.** If G is a definably amenable NIP group, the restriction of the natural map $S_G^{\mathrm{ext}}(M) \to G/G^{00}$ to any ideal group of G is an isomorphism.

Remarkably, the model-theoretic techniques involved in stating, approaching, and proving the conjecture are anything but peculiar to this particular problem, and the main focus of this thesis is on the development and understanding of said techniques. This is reflected in the fact that we will deal with Ellis semigroups only in the first chapter and in the closing section.

We start in Chapter 1 by studying enveloping semigroups, first in the classical context ([Ell69]) and then in the definable one, without any kind of tameness assumption ([New09, New12]). In Chapter 2 we introduce some techniques, still without assuming anything on the underlying theory beyond being first-order complete. In Chapter 3 we introduce dependent theories, see how the previously introduced tools behave in this context, and explore some constructions that heavily exploit the NIP hypothesis. In Chapter 4 we bring in the last ingredient, i.e. definable amenability, see that under our hypotheses it is preserved when passing to Shelah's expansion, characterize it in terms of f-generic types, and conclude by studying the proof of the Ellis Group Conjecture.

A Note on References

We often give references to results as in this example[†].

Theorem ([CK90, Theorem 6.1.15 (Keisler-Shelah)]). Two L-structures are elementarily equivalent if and only if they have isomorphic ultrapowers.

When we do this, we mean that a good place to start searching for an account of the Theorem is Theorem 6.1.15 in [CK90], and that the Theorem is sometimes called "Keisler-Shelah Theorem". This does *not* necessarily imply any of the following, nor it necessarily implies any of their negations.

- The Theorem first appeared in [CK90].
- The statement of the Theorem is exactly as in Theorem 6.1.15 in [CK90].
- The Theorem bears the name of "Keisler-Shelah Theorem" in [CK90].
- The proof we give is the same as in [CK90].

Prerequisites

We take for granted basic notions and results concerning groups, group actions, topological spaces, ultrafilters. From Section 1.2 onwards we also assume some knowledge of model theory, and from Section 2.4 we also assume a little background in measure theory and probability. Also, the reader is expected to be familiar with the basics of set theory, including cardinal arithmetic and cofinality, and with a number of elementary facts such as the triangle inequality. Some of the relevant definitions and theorems are recalled in the Appendix.

Apart from this, the thesis is intended to be self-contained, the only exception being that every once in a while *stable theories* are mentioned in examples, comments, or to provide motivation. We do not assume any knowledge of stable theories, and even those parts are written with this in mind; in any case, they can safely be skipped.

Conventions and Notations

We now summarize some of the conventions, notations and abuses thereof that will be applied throughout this thesis unless explicitly stated otherwise.

All theories and types will be supposed first-order with equality, consistent and complete. Theories are moreover assumed to have infinite models. The first order language in which we work will often be called L, and $|L|$

[†]We have intentionally chosen a theorem that will not appear later on in the thesis.

refers to the cardinality of the set of its *formulas*. For instance, even if $L = \emptyset$, we still have $|L| = \aleph_0$. We usually work in an L-theory T.

Tuples of variables (sometimes infinite ones) are often treated as variables, e.g. we will write $\varphi(x)$ instead of $\varphi(x_0, \dots, x_{n-1})$, $\varphi(\bar{x})$, $\varphi(\vec{x})$, or the like, unless it is necessary to focus the attention on particular coordinates. Also, sometimes we write $\bar{x} = (x_0, \dots, x_{n-1})$ and each x_i is a tuple or variables. The length of the tuple x will be written as $|x|$. Tuples of variables are often denoted with lowercase letters from the end of the latin alphabet.

The same convention applies to tuples of parameters, apart from the fact that they are generally denoted with lowercase letters that appear early in the Latin alphabet; if such letters are uppercase, they generally denote parameter sets.

Models, which are often called M or N, with various sub and superscripts, are notationally identified with their underlying set and $|M|$ will denote its cardinality. Even if $|a| > 1$ we will freely write $a \in M$ instead of $a \in M^{|a|}$. The same applies to definable subsets, i.e. we will say that "$\varphi(x)$ defines a subset of M" even if, strictly speaking, it defines a subset of $M^{|x|}$.

If some variables in $\varphi(x)$ are dummy, we may also mean that $\varphi(x)$ defines a subset of M^k for some $k < |x|$; anyway if this happens it should be clear from context. Definable subsets are often identified with formulas, so we will write things like $\varphi(x) \subseteq A$. Sometimes we simply say φ instead of $\varphi(x)$, and the same applies to types.

Formulas may contain parameters without prior notice; if we want to emphasize that they do not, we will write something like $\varphi(x) \in L$. The set of L-formulas with parameters from A is denoted $L(A)$.

If a set is defined by an $L(A)$-formula we say it is A-definable; the Boolean algebra of all A-definable sets in the variables x is called $\mathrm{Def}_x(A)$. The notation $S_x(A)$ refers to its Stone space[‡], which is identified with the space of types in $|x|$ variables with free variables x with parameters from A. If $\varphi(x)$ (resp. $\pi(x)$) is a formula (resp. partial type) with parameters from A, the corresponding clopen (resp. closed) set will be denoted with $[\varphi(x)]$ (resp. $[\pi(x)]$).

$S(A)$ denotes the space of types in any number of variables. Since we allowed tuples to be infinite, strictly speaking this means that $S(A)$ is a proper class. In practice we write $p \in S(A)$ when it is irrelevant to know who the free variables in p are, but the concerned reader may assume $S(A)$ to be $\bigcup_{n<\omega} S_{x_0,\dots,x_n}(A)$, or $\bigcup_{\alpha<\kappa} S_{(x_i)_{i<\alpha}}(A)$ for some sufficiently large cardinal κ. The same convention applies to $\mathrm{Def}(A)$.

As in [Sim15], the semicolon is used to emphasize a distinction between "object variables" and "parameter variables" inside a formula; for instance, if we write $\varphi(x; y)$ we are probably about to find a parameter b and consider $\varphi(x; b)$ as a definable subset of M. Variables will also happen to "change

[‡]See Definition A.18.

status", for instance we will write something like $\varphi(m; y)$. Sometimes we use $\varphi(x; b) \in L(A)$ as a shorthand for "$\varphi(x; y) \in L$ and $b \in A$".

The notation $a \equiv_A b$ means $\mathrm{tp}(a/A) = \mathrm{tp}(b/A)$. We denote $\mathrm{Aut}(M/A)$ the group of automorphisms f of M that fix A point-wise, i.e. such that $\forall a \in A\ f(a) = a$. We denote the symmetric difference $\neg(\varphi \leftrightarrow \psi)$ with $\varphi \bigtriangleup \psi$. If κ is a cardinal, κ^+ is its successor. The "Standard Lemma" is Lemma A.48.

The Monster Model is denoted $\mathfrak{U}$. A type over the monster will be called a *global* type. See Section A.6 for monster conventions.

Chapter 1

Enveloping Semigroups and Stone Spaces

In this chapter we present the problem whose solution will be the motivation of the present thesis. We will first study a topic from topological dynamics, namely the construction of the *enveloping semigroup of a G-flow*, due to Ellis. Subsequently we will see how Newelski adapted this construction to make it work in a first-order definable setting. At the end of the chapter we will be able to state the aforementioned problem, namely the *Ellis Group Conjecture*.

1.1 Classical Ellis Theory

This section is devoted to the exploration of the ideas of Ellis, as presented in [Ell69]. We will construct the *enveloping semigroup* of a compact G-flow, study its basic properties, and establish the existence of the *Ellis group* associated to the flow. Our exposition follows closely the original one, with the following three notable exceptions. The first one concerns notations and terminology, since we adopt conventions which differ from the original ones; in particular we will write group actions on the *left* and speak of *flows* instead of *transformation groups*. Furthermore, we will assume our topological spaces to be compact; the basic definitions and results (or slight variations of them) are still valid for locally compact Hausdorff spaces, but since compactness is needed in order to define the enveloping semigroup we prefer to assume it from the beginning; as a by-product, some proofs and statements are simplified. Finally, in order to foreshadow the adaptation to the definable context, we will use filters instead of nets; this facilitates the translation process due to the fact that (ultra)filters have a natural definable analogue, namely (complete) types. See Section A.3 for basic results on the use of filters in topology.

G-Flows

Definition 1.1. Let $(G,\cdot)$ be a group. A *G-flow* is a continuous action of G on a compact Hausdorff space.

With the words "continuous action" on a topological space X we mean that G is a topological group and the action is continuous as a map $G\times X\to X$. Anyway, unless mentioned otherwise, we will put on G the discrete topology; hence G will automatically be a topological group and a continuous action will simply be a group homomorphism $\vartheta\colon G\to \mathrm{Homeo}(X,X)$. For $g\in G$ and $p\in X$, the image $(\vartheta(g))(p)$ of p under the homeomorphism corresponding to g will be also denoted in one of the following ways

$$\vartheta_g(p)=g\cdot p=gp=\alpha_p(g)$$

We usually identify a G-flow with the underlying topological space X (e.g. "a G-flow X") whenever the action is unspecified or clear from context.

Example 1.2. Any homeomorphism f of a compact Hausdorff space X into itself gives rise to a $\mathbb{Z}$-flow on X with the action given by $n\cdot p=f^n(p)$.

Definition 1.3. A *morphism of G-flows* is a continuous $\varphi\colon X_0\to X_1$ that commutes with the action of g, i.e. such that for all $p\in X_0$ and $g\in G$ we have $\varphi(g\cdot p)=g\cdot\varphi(p)$.

Example 1.4. If (X_0,f_0) and (X_1,f_1) are two $\mathbb{Z}$-flows as in Example 1.2, a morphism $X\to Y$ is a continuous map $\varphi\colon X\to Y$ such that $f_1\circ\varphi=\varphi\circ f_0$.

Definition 1.5. A *subflow* of a G-flow X is a *non-empty*, *closed*, *G-invariant* subset of X. In symbols, $\emptyset\neq G\cdot Y\subseteq Y=\overline{Y}\subseteq X$.

Example 1.6. The *orbit closure* $\overline{Gp}$ of a point $p\in X$ is a subflow of X: it is obviously a closed subspace, and it is closed under the action of G because each ϑ_g is a continuous map.

Definition 1.7. A *minimal subflow* is a subflow which is minimal under inclusion.

Proposition 1.8. Minimal subflows exist.

Proof. Every descending chain of subflows is still G-invariant and has non-empty intersection by compactness. Since X is itself a subflow, we can apply Zorn's Lemma to the family of subflows ordered by reverse inclusion and the proposition follows. □

Definition 1.9. A subset $A\subseteq G$ will be called *syndetic* if there is a finite $F\subseteq G$ such that $G=F\cdot A$.

If G is a topological group, not necessarily equipped with the discrete topology, a more general notion of "syndetic" is obtained replacing "finite" with "compact" in the previous definition. As we said before, our groups carry the discrete topology unless mentioned otherwise, so for our purposes the two definitions are equivalent.

Definition 1.10. A point $p \in X$ is *almost periodic* iff, for each of its neighbourhoods U, the set $\alpha_p^{-1}(U) = \{g \in G \mid gp \in U\}$ is syndetic in G.

An etymological explanation of the diction "almost periodic", together with a motivation for the introduction of this notion, can be found at the beginning of [Ell69, Chapter 2].

The following proposition is a special case of [Ell69, Proposition 2.5].

Proposition 1.11. A point p is almost periodic if and only if its orbit closure is a minimal subflow.

Proof. ($\Rightarrow$) It is enough to show that, for every $q \in \overline{Gp}$, we have $p \in \overline{Gq}$. Hence, given a compact[1] neighbourhood U of p, we need to prove that $Gq \cap U \neq \emptyset$. Since, by hypothesis, there are $g_0, \ldots, g_{n-1}$ such that $G = \bigcup_{i<n} g_i \cdot \alpha_p^{-1}(U)$, we have

$$q \in \overline{Gp} = \overline{\bigcup_{i<n} g_i \alpha_p^{-1}(U) \cdot p} \subseteq \overline{\bigcup_{i<n} g_i U} = \bigcup_{i<n} g_i U$$

where the last equality follows from the fact that the compact subspaces of an Hausdorff space are closed. This implies that $g_i^{-1} q \in U$ for some $i < n$.

($\Leftarrow$) If U is an open neighbourhood of p then $\overline{Gp} \setminus GU$ is a closed, invariant, and proper subset of $\overline{Gp}$, hence it is empty and by compactness there are finitely many $g_i \in G$ such that $Gp \subseteq \overline{Gp} \subseteq \bigcup_{i<n} g_i U$. But this means that for every $g \in G$ there is $i < n$ such that $gp \in g_i U$, i.e. $g_i^{-1} g \in \alpha_p^{-1}(U)$ and thus $G = \bigcup_{i<n} g_i \alpha_p^{-1}(U)$. □

Example 1.12 ([New09, Example 1])**.** Consider the action of $(\mathbb{Z}, +)$ by shift on the space $2^{\mathbb{Z}}$ with the usual product topology. If we think of elements of $2^{\mathbb{Z}}$ as infinite (on both sides) strings, then it is easy to see that the orbit of any concatenation η of all finite strings[2] is dense, and a proper subflow can be found simply by considering one of the two constant strings. Hence η is *not* almost periodic.

Remark 1.13. The previous equivalence need not be true if one is working with a semigroup action. For instance, if we let $(\mathbb{Z}, \cdot)$ (or even $(\mathbb{Z} \setminus \{0\}, \cdot)$) act on $\mathbb{R}/\mathbb{Z}$, then the equivalence class of $1/2$ is almost periodic, but its orbit (which is already closed) contains the fixed point corresponding to the equivalence class of 0.

[1] See Proposition A.3.

[2] To obtain such an η fix a bijection $f\colon \mathbb{Z} \to 2^{<\omega}$ and, starting with $f(0)$, recursively concatenate $f(n)$ to the right and $f(-n)$ to the left.

The Enveloping Semigroup

Definition 1.14. If X is a G-flow, its *enveloping semigroup* $E(X)$ is the closure of $\{\vartheta_g \mid g \in G\}$ with respect to the product topology on X^X.

We will now study the basic properties of $E(X)$ and explain the "semigroup" diction. First, we recall the notion of limit of a filter. Some basic theorems and further definitions concerning filters and limits are collected in Section A.3.

Definition 1.15. If $\mathcal{F}$ is a filter on a topological space Y, and $\ell \in Y$, we say that $\ell = \lim \mathcal{F}$ iff every neighbourhood of ℓ is in $\mathcal{F}$. If $f \colon Z \to Y$ is a function[3] from a set Z to Y and $\mathcal{F}$ is the pushforward $f^*(\mathcal{F}_0)$ of a filter $\mathcal{F}_0$ on Z, we also write $\lim_{z \to \mathcal{F}_0} f(z)$ for $\lim \mathcal{F}$.

Lemma 1.16. *Every element of $E(X)$ can be written as $\vartheta_{\mathcal{U}} = \lim_{g \to \mathcal{U}} \vartheta_g$ for a suitable (not necessarily unique) ultrafilter $\mathcal{U}$ on G. Conversely, for every ultrafilter $\mathcal{U}$ on G, we have $\vartheta_{\mathcal{U}} \in E(X)$.*

Proof. If $p \in E(X)$ each of its neighbourhoods U contains some ϑ_g by definition. Then the family $\{\{g \in G \mid \vartheta_g \in U\} \mid U \text{ a neighbourhood of } p\}$ extends to an ultrafilter $\mathcal{U}$, and $\lim_{g \to \mathcal{U}} \vartheta_g = p$ by construction. The second statement follows from the fact that $E(X)$ is closed. □

Notation 1.17. As usual, we will identify every $g \in G$ with the corresponding principal ultrafilter $\sqcup_g$. Obviously $\vartheta_g = \vartheta_{\sqcup_g}$, so no confusion arises. For this reason, we will use the same notation to refer to the action of G on $E(X)$ and, after Proposition 1.22, to the semigroup operation on $E(X)$.

Proposition 1.18. *$E(X)$ is a G-flow with the action given by $g \cdot \vartheta_{\mathcal{U}} = \vartheta_g \circ \vartheta_{\mathcal{U}}$.*

Proof. $E(X)$ is a closed subspace of X^X, which is compact Hausdorff because X is by Tychonoff's theorem, hence it is compact Hausdorff itself. Set[4] $g\mathcal{U} = g \cdot \mathcal{U} = (g \cdot -)^*(\mathcal{U})$. It is easy to check that $\vartheta_g \circ \vartheta_{\mathcal{U}} = \vartheta_{g\mathcal{U}} \in E(X)$. Since ϑ_g is continuous and the topology on $E(X)$ is induced by the product topology on X^X, i.e. the topology of point-wise convergence, it follows that $\vartheta_g \circ - \colon E(X) \to E(X)$ is continuous, and it obviously has $\vartheta_{g^{-1}} \circ -$ as an inverse. Moreover $(gh) \cdot \vartheta_{\mathcal{U}} = \vartheta_{gh} \circ \vartheta_{\mathcal{U}} = \vartheta_g \circ \vartheta_h \circ \vartheta_{\mathcal{U}} = g \cdot (h \cdot \vartheta_{\mathcal{U}})$. □

Recall the following definition:

Definition 1.19. If $\mathcal{U} \in \beta X$ and $\mathcal{V} \in \beta Y$, the *tensor product* $\mathcal{U} \otimes \mathcal{V} \in \beta(X \times Y)$ is defined as

$$U \in \mathcal{U} \otimes \mathcal{V} \iff \{x \in X \mid \{y \in Y \mid (x, y) \in U\} \in \mathcal{V}\} \in \mathcal{U}$$

[3] It may be useful to think of the special case where f is a sequence, i.e. a function $\omega \to Y$.

[4] This $g \cdot -$ is the multiplication map $G \to G$ than sends h to gh. It is *not* the action map of g on X.

Definition 1.20. If $\mathcal{U},\mathcal{V} \in \beta G$, then $\mathcal{U}\cdot\mathcal{V}$ is the pushforward of $\mathcal{U}\otimes\mathcal{V}$ under the multiplication map $G \times G \to G$, i.e.

$$U \in \mathcal{U}\cdot\mathcal{V} \iff \{g \in G \mid \{h \in G \mid gh \in U\} \in \mathcal{V}\} \in \mathcal{U} \tag{1.1}$$

Lemma 1.21. The operation $- \cdot -\colon \beta G \times \beta G \to \beta G$ is associative.

Proof. Being associative is inherited from the group product of G, as if we try to check if $U \in (\mathcal{U}\cdot\mathcal{V})\cdot\mathcal{W}$ and if $U \in \mathcal{U}\cdot(\mathcal{V}\cdot\mathcal{W})$, in both cases we end up checking if $\{g \mid \{h \mid \{\ell \mid g\cdot h\cdot \ell \in U\} \in \mathcal{W}\} \in \mathcal{V}\} \in \mathcal{U}$. □

So this puts a semigroup operation on βG. This does not happen by chance, and we will come back on said semigroup structure later on.

Proposition 1.22. $E(X)$ is a semigroup with the operation of composition. Namely, $\vartheta_{\mathcal{U}} \circ \vartheta_{\mathcal{V}} = \vartheta_{\mathcal{U}\cdot\mathcal{V}}$.

Proof. By continuity of ϑ_g and the fact that we are using point-wise limits we have $\vartheta_g \circ \lim_{h\to\mathcal{V}} \vartheta_h = \lim_{h\to\mathcal{V}} \vartheta_g \circ \vartheta_h$, hence

$$\vartheta_{\mathcal{U}} \circ \vartheta_{\mathcal{V}} = \lim_{g\to\mathcal{U}} \lim_{h\to\mathcal{V}} \vartheta_g \circ \vartheta_h = \lim_{g\to\mathcal{U}} \lim_{h\to\mathcal{V}} \vartheta_{g\cdot h} = \lim_{(g,h)\to\mathcal{U}\otimes\mathcal{V}} \vartheta_{g\cdot h} = \lim_{\ell\to\mathcal{U}\cdot\mathcal{V}} \vartheta_\ell$$

See Lemma A.31 for more details on the last two equalities. □

Corollary 1.23. If $\vartheta_{\mathcal{U}_0} = \vartheta_{\mathcal{U}_1}$ and $\vartheta_{\mathcal{V}_0} = \vartheta_{\mathcal{V}_1}$ then $\vartheta_{\mathcal{U}_0\cdot\mathcal{V}_0} = \vartheta_{\mathcal{U}_1\cdot\mathcal{V}_1}$

Proof. They both are equal to $\vartheta_{\mathcal{U}_0} \circ \vartheta_{\mathcal{V}_0} = \vartheta_{\mathcal{U}_1} \circ \vartheta_{\mathcal{V}_1}$. □

The reader may be puzzled by the fact that we call $E(X)$ the enveloping *semigroup* instead of enveloping *monoid*: after all $\mathrm{id}_X \in E(X)$. The reason behind this is the following: we will soon be interested in certain sub-semigroups of $E(X)$ that will happen to be groups under composition *but with a group identity different from* id_X. Calling $E(X)$ the "enveloping monoid" would suggest that its substructures have to contain id_X and exclude the aforementioned groups. In more model-theoretic terms, we will be interested in substructures of $(E(X), \circ)$, which need not be substructures of $(E(X), \circ, \mathrm{id}_X)$.

Proposition 1.24. The map $-\circ\vartheta_{\mathcal{U}}\colon E(X) \to E(X)$ is continuous. If $g \in G$, then also $\vartheta_g \circ -\colon E(X) \to E(X)$ is continuous.

Proof. This follows easily from the fact that the topology on $E(X)$ is the one of point-wise convergence and the ϑ_g are continuous maps. □

Remark 1.25. The maps $\vartheta_{\mathcal{U}} \circ -\colon E(X) \to E(X)$ need not be continuous in general. A moment's thought reveals that this would be equivalent to the $\vartheta_{\mathcal{U}}\colon X \to X$ being continuous, and this need not be true, since continuity is not generally preserved under point-wise limits.

Ellis Semigroups

The following definition encompasses the core properties of $E(X)$.

Definition 1.26. An *Ellis semigroup* is a compact Hausdorff space equipped with a semigroup operation which is continuous in the first coordinate, i.e. for all p the map $- \cdot p$ is continuous.

The following result is essential to the development of the theory. Apart from having a slightly stronger hypothesis[5] it is essentially [Ell69, Lemma 2.9].

Theorem 1.27. Ellis semigroups always have idempotents.

Proof. Given an Ellis semigroup $(E, \cdot)$, define the following family

$$\mathcal{C} = \{C \subseteq E \mid C = \overline{C} \neq \emptyset,\ C \cdot C \subseteq C\}$$

Clearly $E \in \mathcal{C}$ and, by compactness, $\mathcal{C}$ satisfies the hypotheses of Zorn's Lemma when ordered by reverse inclusion, so it has a minimal element $\tilde{C}$. Let $x \in \tilde{C}$. We will show that $x = x^2$, and indeed $\tilde{C} = \{x\}$.

Claim. $\tilde{C} \cdot x \in \mathcal{C}$

Proof of the Claim. Since $-\cdot x$ is continuous, $\tilde{C}\cdot x$ is the image of a continuous map from a compact space to an Hausdorff one, hence it is closed. Moreover

$$(\tilde{C} \cdot x) \cdot (\tilde{C} \cdot x) = \tilde{C} \cdot (x \cdot \tilde{C}) \cdot x \subseteq \tilde{C} \cdot \tilde{C} \cdot x = (\tilde{C} \cdot \tilde{C}) \cdot x \subseteq \tilde{C} \cdot x$$

□ CLAIM

Since $\tilde{C} \cdot x \subseteq \tilde{C} \cdot \tilde{C} \subseteq \tilde{C}$, by minimality $\tilde{C} \cdot x = \tilde{C}$. This proves that the following set is non-empty:

$$D = \{y \in \tilde{C} \mid y \cdot x = x\} = (- \cdot x)^{-1}(x) \cap \tilde{C}$$

Clearly D is a closed set. Moreover $D \cdot D \subseteq D$, because if $y_0, y_1 \in D$

$$(y_0 \cdot y_1) \cdot x = y_0 \cdot (y_1 \cdot x) = y_0 \cdot x = x$$

Hence $D \in \mathcal{C}$ and, since $D \subseteq \tilde{C}$ by definition, by minimality $D = \tilde{C}$. This means that $x \in D$, i.e. $x \cdot x = x$. □

Definition 1.28. A *left ideal* of a semigroup $(S, \cdot)$ is a *non-empty* $I \subseteq S$ which is left-absorbing, i.e. such that $S \cdot I \subseteq I$. When we simply say that I is an *ideal* we mean that I is a left ideal.

The interest in ideals lies, to begin with, in the following results.

Proposition 1.29 ([Ell69, Proposition 3.4])**.** In the enveloping semigroup minimal subflows coincide with minimal ideals.

[5] Namely, in order for the present proof to be carried out, it suffices to require E to be compact T1 and the map $- \cdot p$ to be continuous *and closed*.

Proof. Let M be a minimal subflow of the G-flow $E(X)$ and let $p \in M$. Since M is G-invariant $Gp \subseteq M$, and since M is closed $E(X)p = \overline{Gp} \subseteq \overline{M} = M$. This proves that M is an ideal. If $I \subseteq M$ is another, given $p \in I$ we can write

$$\overline{Gp} = E(X)p \subseteq E(X)I \subseteq I \subseteq M$$

Since M is a minimal subflow all the inclusions above are equalities, and M is thus a minimal ideal.

If now I is a minimal ideal, then it is G-invariant because $GI \subseteq E(X)I \subseteq I$. For any $p \in I$, we have that $E(X)p$ is clearly an ideal contained in I, and by minimality $E(X)p = I$. This shows simultaneously that I is closed (because $- \cdot p$ is continuous and $E(X)$ is compact), and that I is minimal as a subflow, since $\overline{Gp} = E(X)p = I$ and $p \in I$ was arbitrary. □

Anyway the existence of minimal ideals and the fact that they are closed is not peculiar to envelopes, as the following to results show.

Proposition 1.30. Minimal ideals of an Ellis semigroup are closed.

Proof. As in the proof of Proposition 1.29, it suffices to write a minimal ideal I of the semigroup E as $I = Ep$, where p is any element of I. □

Proposition 1.31. Every ideal of an Ellis semigroup includes a minimal ideal. In particular minimal ideals always exist.

Proof. Given an ideal I of E, apply Zorn's Lemma to the family of *closed* ideals $J \subseteq I$ ordered by reverse inclusion, which is non-empty because for all $p \in I$, by continuity, Ep belongs to the family. □

The following theorem sums up some "working facts" concerning Ellis semigroups. Its proof, despite its length, is elementary.

Theorem 1.32 ([Ell69, Propositions 3.5 and 3.6 (1)])**.** Let I be a minimal ideal of an Ellis semigroup E. Then the following facts hold:

1. The set $\mathrm{Idem}(I)$ of idempotents of I is non-empty.
2. If $p \in I$ and $u \in \mathrm{Idem}(I)$, then $pu = p$.
3. If $u \in \mathrm{Idem}(I)$, then uI is a subgroup of I with identity u.
4. $\{uI \mid u \in \mathrm{Idem}(I)\}$ is a partition of I.
5. If I, J are minimal ideals and $u \in \mathrm{Idem}(I)$, then there is a unique $v \in \mathrm{Idem}(J)$ such that $vu = u$ and $uv = v$.

Proof.

1. A minimal ideal, being closed, is an Ellis semigroup itself. It suffices then to apply Theorem 1.27 to I.

2. Since Iu is an ideal contained in I, by minimality they coincide, and therefore there is $q \in I$ such that $p = qu$. Then $pu = quu = qu = p$.

3. Writing $uIuI = u(IuI)$ and observing that IuI is an ideal included in I yields, by minimality, that uI is closed under the semigroup operation, that trivially remains associative.

   If $up \in uI$, a fortiori $up \in I$, so by point 2 u is a right identity for up. On the other hand $uup = up$ by idempotence, hence u is also a left identity for all elements of uI.

   We are left to find and inverse for $q = up$. Since $Iq = I$ there is $r \in I$ such that $rq = u$. We claim that ur is an inverse for q in uI. Indeed, multiplying on the left yields $urq = uu = u$, hence every element of uI has a left inverse in uI. Applying the last statement to ur yields an $s \in uI$ such that $sur = u$. Then $q = uq = surq = su = s$. But then $u = sur = qur$ and ur is a right inverse of q.

4. Fix any $p \in I$. By minimality $Ip = I$, and this means that $\{q \in I \mid qp = p\}$ is non-empty. An easy check shows that it is a sub-semigroup of I, and it is moreover closed since it can be written as $I \cap (- \cdot p)^{-1}(p)$. Then Theorem 1.27 applies and yields $u \in \mathrm{Idem}(I)$ such that $up = p$, witnessing $p \in uI$. This shows that $I = \bigcup_{u \in \mathrm{Idem}\, I} uI$, but we still have to prove that this union is disjoint. Suppose $u, v \in \mathrm{Idem}(I)$ and let $p \in uI \cap vI$. Let q be the inverse of p in vI. Taking advantage of point 2 and of the fact that $up = p$ (because $p \in uI$) we have $u = uv = upq = pq = v$.

5. Ju is an ideal and equals I by minimality. As in the previous point, this allows us to apply Theorem 1.27 to $\{p \in J \mid pu = p\}$ to find $v \in \mathrm{Idem}(J)$ such that $vu = u$. Applying the same reasoning to v yields $w \in \mathrm{Idem}(I)$ such that $wv = v$. Since $w \in I$, by point 2 $wu = w$. But then $w = wu = wvu = vu = u$. As for uniqueness: if $v, w \in \mathrm{Idem}(J)$ both satisfy the thesis, then $v = uv = wuv = wv = w$, where the last equality is again by point 2. □

Remark 1.33. We have partitioned each minimal ideal in groups. Since minimal ideals are disjoint[6] we could wonder whether they partition E, i.e. whether every point is contained in a minimal ideal. This is not true in general, and it is easy to see that if E is the enveloping semigroup of a G-flow it is equivalent to every point of E being almost periodic.

Definition 1.34. If E is an Ellis semigroup, an *Ellis group* or *ideal group* of E is one of its subgroups of the form uI for I a minimal ideal and $u \in \mathrm{Idem}(I)$.

[6]If the intersection of two ideals is non-empty, it is still an ideal.

Theorem 1.35. All the Ellis groups of an Ellis semigroup are isomorphic.

Proof. We show first that all the groups inside a given minimal ideal I are isomorphic. Indeed consider the maps $u_0 I \xrightarrow{u_1 \cdot -} u_1 I$ and $u_1 I \xrightarrow{u_0 \cdot -} u_0 I$.

- By point 2 of Theorem 1.32, for $i \in \{0, 1\}$ and $p, q \in u_{1-i} I \subseteq I$ we have $u_i p u_i q = u_i pq$, so our maps are homomorphisms.
- For the same reason and the fact that the $u_i I$ are groups, if $p \in u_i I$ we have $u_i u_{1-i} p = u_i p = p$. Hence the two maps are homomorphisms and inverses of each other.

Now we show that, for I, J minimal ideals and $u \in \mathrm{Idem}(I)$, every uI has a "twin" inside J. Let v be given by point 5 of Theorem 1.32. Then, if $p = vq \in vJ$, we have $pu = vqu = uvqu = upu \in uI$. We can then consider the map $vJ \xrightarrow{- \cdot u} uI$ and, similarly, the map $uI \xrightarrow{- \cdot v} vJ$.

- Let $p_0, p_1 \in vJ$ as witnessed by writing $p_i = vq_i$. Since $uv = v$,

$$p_0 u p_1 u = p_0 u v q_1 u = p_0 v q_1 u = p_0 p_1 u$$

  so this map is an homomorphism. Similarly for the other one.
- If $p \in vJ$ we have $puv = pv = p$, and similarly if $q \in uI$ we have $qvu = qu = q$. Hence the two maps are inverses of each other.

Therefore, for all uI and wJ we can find an isomorphism composing the two above: namely, if v is given by point 2 of Theorem 1.32 applied to u, the map $uI \to wJ$ given by $w \cdot - \cdot v$ is an isomorphism, with inverse $v \cdot - \cdot u$. The skeptic reader may be convinced by the following (redundant) verification:

$$vwpvu = vwpu = vwp = vp = vu\tilde{p} = u\tilde{p} = p$$
$$wvquv = wvqv = wvq = wq = q$$

□

We are therefore authorized to speak of *the* Ellis group associated to an Ellis semigroup.

Definition 1.36. The *Ellis group* (or the *ideal group*) of a G-flow X is the Ellis group associated to $E(X)$.

Another beautiful consequence of Theorem 1.32 is the following result.

Theorem 1.37 ([Ell69, Proposition 3.6 (3)[7]])**.** Minimal subflows of $E(X)$ are all isomorphic as G-flows.

Proof. Let I, J be two minimal subflows of $E(X)$. By Proposition 1.29 they are minimal ideals. Fix $u \in \mathrm{Idem}\, I$ and let $v \in \mathrm{Idem}\, J$ be given by point 5 of Theorem 1.32. Then $- \cdot v \colon I \to J$ and $- \cdot u \colon J \to I$ are inverses of each other, we already know that they are continuous, and the fact that they commute with the action of G is trivial. □

[7] Probably due to a misprint, in the first edition of [Ell69] (or at least in the copy available in the library of the University of Pisa), the thesis of Proposition 3.6 has no point (3) (but the proof does).

The Universal G-Ambit

We now turn our attention to a very special G-flow: the space βG of ultrafilters on G, with the action of G given by[8] $g \cdot \mathcal{U} = g\mathcal{U} = (g \cdot -)^*(\mathcal{U})$. If the reader is familiar with Stone-Čech compactifications, she will probably now expect βG to enjoy some universal property as a G-flow. This intuition is correct, as explained in [Ell69, Chapter 7], of which the following paragraph is intended to be a brief summary.

A G-*ambit* is a G-flow together with a distinguished point whose orbit is dense. For instance, every enveloping semigroup $E(X)$ is a G-ambit, since the orbit closure of ϑ_e, where e is the identity of G, is the whole $E(X)$ by definition. The G-flow βG is a G-ambit too, since it can be easily checked that the principal ultrafilter $\sqcup_e$ on the group identity e has dense orbit. Thus, if X is a G-ambit with distinguished point p, one shows using the properties of the Stone-Čech compactification that the map sending $\sqcup_g$ to gp extends uniquely to a surjective morphism of G-ambits $(\beta G, \sqcup_e) \to (X, p)$. Thus βG is the "largest" G-ambit, in the sense that each G-ambit is an homomorphic image of βG. One then proves that every surjective morphism of G-ambits $X \to \beta G$ is indeed an isomorphism, and exploits this fact to show that $E(\beta G)$ and βG are isomorphic, the isomorphism simply being evaluation in $\sqcup_e$. That is not the end of the story: minimal subflows of βG, in addition to being all isomorphic (by Theorem 1.37 and the previous isomorphism), share a similar universal property among minimal G-flows. For a more thorough discussion we refer the reader to [Ell69] and to the monograph [Aus88].

For our purposes it is better to write down the isomorphism between βG and $E(\beta G)$ explicitly, since it will shed light on the constructions of the following section.

Lemma 1.38. Every $f \in E(\beta G)$ is of the form $\mathcal{U} \cdot -$ for a suitable $\mathcal{U} \in \beta G$.

Proof. By Lemma 1.16 there is $\mathcal{U} \in \beta G$ such that $f = \vartheta_{\mathcal{U}} = \lim_{g \to \mathcal{U}} \vartheta_g$. We claim that such an $\mathcal{U}$ serves our purposes. In order to prove this, given $\mathcal{V} \in \beta G$, we have to show that $\mathcal{U} \cdot \mathcal{V} = f(\mathcal{V})$. For all $U \subseteq G$ we have

$$\{g \mid \{h \mid gh \in U\} \in \mathcal{V}\} = \{g \mid g^{-1}U \in \mathcal{V}\} = \{g \mid U \in g\mathcal{V}\} = \{g \mid U \in \vartheta_g(\mathcal{V})\}$$

Now it suffices to notice that the leftmost set belongs to $\mathcal{U}$ if and only if $U \in \mathcal{U} \cdot \mathcal{V}$, and the rightmost set belongs to $\mathcal{U}$ if and only if $U \in f(\mathcal{V}) = \Big(\lim_{g \to \mathcal{U}} \vartheta_g\Big)(\mathcal{V}) = \lim_{g \to \mathcal{U}} \vartheta_g(\mathcal{V})$. □

Theorem 1.39 ([Ell69, Corollary 7.12])**.** $E(\beta G)$ and βG are isomorphic G-flows.

[8] The leftmost dot is the symbol for the group action, the rightmost one is the symbol for group multiplication in G. Due to the fact that, when one identifies $g \in G$ with $\sqcup_g \in \beta G$ then we have $g \cdot \sqcup_h = \sqcup_{g \cdot h}$ we feel that this rather standard abuse of notation is justified.

Proof. Consider the map $\vartheta\colon \beta G \to E(\beta G)$ that sends $\mathcal{U}$ to[9] $\vartheta_{\mathcal{U}} = \mathcal{U} \cdot -$. Its inverse will be the map $\vartheta^{-1}\colon E(\beta G) \to \beta G$ sending f to $f(\sqcup_e)$, as it is easy to check given the previous lemma, and it is not difficult to verify that they commute with the action of G, so we only have to take care of continuity issues. Since the map ϑ^{-1} is nothing else than a projection, it is continuous by definition of product topology. If we show that it is also open we are done, so we fix a basic open set of $E(\beta G)$, which must be of the form $\{f \mid \forall i < n\ f(\mathcal{V}_i) \in [V_i]\} = \{f \mid \forall i < n\ V_i \in f(\mathcal{V}_i)\}$ and we want to check that $\{\mathcal{U} \mid \forall i < n\ V_i \in \mathcal{U} \cdot \mathcal{V}_i\}$ is open in βG. This is obvious since it can be written as $\bigcap_{i<n}[\{g \mid g^{-1}V_i \in \mathcal{V}_i\}]$. □

Corollary 1.40. $(\beta G, \cdot)$ is an Ellis semigroup.

Proof. We could check continuity in the first coordinate directly, but it suffices to observe that by Proposition 1.22 the operation $\mathcal{U} \cdot \mathcal{V}$ is precisely the image of the composition in $E(\beta G)$ under the identification given by the previous theorem. □

While the previous constructions and ideas are central to the present thesis, we need to concentrate our efforts on their definable counterparts, and are therefore forced to omit a large number of their consequences and related notions such as the concepts of *distal* or *proximal* systems. Nonetheless the study of enveloping semigroups has been very fruitful in topological dynamics. The interested reader can find more on these topics in — for instance — [Ell69, Gla07a, Gla07b, AAG08].

Even if we start with a semigroup S instead of a group G some of the constructions presented above still work, and endow βS with an Ellis semigroup structure. This fact has found applications in additive combinatorics and Ramsey theory, for instance giving an extremely elegant proof of Hindman's Theorem that exploits minimal ideals and idempotents of $(\beta\mathbb{N}, +)$, or proving in a similar fashion the classical result of Van der Waerden on monochromatic arithmetic progressions. For an account, see [HS98]. The interaction between dynamics, algebra and combinatorics is also studied in [Bla93].

1.2 Newelski's Set-Up

In [New09] Newelski "imported" topological dynamics into model theory in order to generalize the concept of *generic type* from stable group theory (see [Poi01, Pil96]) to an arbitrary first-order theory and, using *coheirs*, he was able to adapt Ellis theory to the definable case. The study of "definable Ellis theory" was then continued in [New12], where the topic was approached

[9]This is not an abuse of notation: with the current action $\vartheta_{\mathcal{U}}$ is precisely the same as the one defined in Lemma 1.16; the skeptic reader will probably agree after writing down the definitions.

from the different angle of *externally definable subsets*. As we shall soon see these are nothing but two sides of the same coin, the former being more convenient to handle when dealing with computations, the latter probably providing a more conceptually clear framework. Our exposition will be closer to [New12] at the beginning of the section and closer to [New09] at the end.

From now on we assume the reader to be familiar with basics of model theory; some of the more frequently used facts are collected in Section A.5. Conventions and notations are recalled in the Introduction.

Definable Groups

Definition 1.41. Fix an L-theory T and work modulo T. A *definable group* is given by

- a formula φ defining a non-empty subset G,
- a definable function $- \cdot -\colon G \times G \to G$, and
- a definable function $-^{-1}\colon G \to G$

such that G equipped with $\cdot$ and $^{-1}$ satisfies the group axioms. If G can be defined with formulas in $L(A)$, we will say that G is *A-definable* or *definable over A*.

Remark 1.42. A definable group in a theory T can be thought of as the functor that associates to every model M of T the interpretation $G(M)$ of G in M, the morphisms being elementary immersions on the left and injective group homomorphisms on the right.

Example 1.43. In theories of fields (e.g. $\mathsf{ACF}_0 = \mathrm{Th}(\mathbb{C})$ or $\mathsf{RCF} = \mathrm{Th}(\mathbb{R})$) matrix groups such as $\mathrm{GL}(n, -)$, $\mathrm{SL}(n, -)$, $\mathrm{O}(n, -)$, or $\mathrm{SO}(n, -)$, are definable groups, since they (and matrix multiplication and inverse) are defined by polynomial equations.

Notation 1.44. $\varphi(x \cdot y)$ means[10] $\exists z\; x \cdot y = z \wedge \varphi(z)$. Similarly for $\varphi(x^{-1})$.

Definable Types and Externally Definable Sets

Fix an L-theory T, a model M and an $\emptyset$-definable group G, and let $S_G(M)$ be the space of types that concentrate on G, i.e. the closed subset of the space of types with parameters from M in the right number of variables that contain the formula defining G. The key observation is that, since types with parameters from M can be seen as ultrafilters on the Boolean algebra[11]

[10] Remember that a function is definable iff its graph is definable (by definition), hence $x \cdot y = z$ is a shorthand for the formula defining the graph.

[11] See Remark A.38.

of the M-definable subsets of[12] $\mathfrak{U}$, we can hope that the machinery developed in the previous section still works when we replace βG with $S_G(M)$.

To begin with, $S_G(M)$ carries a natural $G(M)$-flow structure given by $g \cdot p(x) = gp(x) = p(g^{-1} \cdot x)$, where $\varphi(x) \in gp(x)$ iff $g^{-1}\varphi(x) := \varphi(g \cdot x) \in p(x)$. We could then expect $S_G(M)$ to carry some of the universal properties enjoyed by βG inside the realm of "definable G-flows". We can for instance start by asking ourselves the more precise question: is $E(S_G(M))$ isomorphic to $S_G(M)$, if possible with "the same" isomorphisms as before? To provide an answer, some tools will be required; before diving into definitions and lemmas, we explain where the need for them comes from.

In order to hope for the question above to have a positive answer, it is necessary for us to be able to equip $S_G(M)$ with an Ellis semigroup structure, since an isomorphism with $E(S_G(M))$ will inevitably provide one. We can then try to "copy and paste" the semigroup operation of Definition 1.20, replace sets with formulas, ultrafilters with types, G with $G(M)$ and then try to make sense of the resulting string. Let us have a look at it:

$$\varphi(x) \in p(x) \cdot q(x) \iff \{g \in G(M) \mid \{h \in G(M) \models \varphi(gh)\} \in q(x)\} \in p(x)$$

The string of symbols above, strictly speaking, does not have any meaning, since types are sets of formulas and it is not clear what it means for a subset of $G(M)$ to belong to $q(x)$. Anyway, looking at a type as an ultrafilter on a Boolean algebra of definable sets eliminates this issue. In other words, the string begins to make more sense if we rewrite it like this:

$$\varphi(x) \in p(x) \cdot q(x) \iff \underbrace{\{g \in G(M) \mid \varphi(g \cdot x) \in q(x)\}}_{\text{?⃝}} \in p(x) \tag{1.2}$$

We have eliminated a problem, but we still have to deal with the set ?⃝. Let us begin by giving it a name.

Definition 1.45. For $q \in S_G(M)$ and $\varphi(x) \in L(M)$ we define

$$\delta_q\varphi = \{g \in G(M) \mid \varphi(g \cdot x) \in q(x)\}$$

Here an obstruction shows up: if p was an ultrafilter on the whole powerset of $G(M)$, asking wheter $\delta_q\varphi \in p$ would be a legitimate question, and if the answer is "no" then p would contain $(\delta_q\varphi)^{\complement}=\delta_q\neg\varphi$ and $p \cdot q$ would be a complete type. But p only cares about M-*definable* subsets of $G(M)$. In other words, the problem is that $\delta_q\varphi$ *need not be* M*-definable.* The following definition is then natural:

[12] Or, if the reader prefers, the Boolean algebra of $L(M)$-formulas modulo equivalence in the elementary diagram of M.

Definition 1.46. A type $p \in S_x(M)$ is *A-definable* iff for all $\varphi(x;y) \in L$ the set $d_{p(x)}\varphi(x;y) = \{b \in M \mid \varphi(x;b) \in p\}$ is A-definable. We say that p is *definable* iff it is A-definable for some $A \subseteq M$, equivalently iff it is M-definable. The set $d_p\varphi$ is called the *defining schema* of p evaluated in φ, and, if p is definable, is often identified with a formula $\psi(y)$ defining it.

If, for some reason, we happen to know than all types in $S_G(M)$ are definable, then (1.2) makes sense and defines an element of $S_G(M)$: if $\varphi(x;b) \in L(M)$ all we need to do is to apply definability of p to the formulas $\psi(x;y,z) = \varphi(y \cdot x;z)$ and then set

$$\delta_{p(x)}\varphi = \big(d_{p(x)}\psi(x;y,z) \wedge z = b\big) \in L(M)$$

This allows us to give the following definition:

Definition 1.47. Suppose that all types over M are definable. We define $p \cdot q \in S_G(M)$ as

$$\varphi \in p \cdot q \iff \delta_q\varphi \in p$$

As the reader may expect, it is *not* always the case that all types in $S_x(M)$ are definable[13]. In order to deal with this, the following observation is now crucial:

Remark 1.48. If p is realized in M by a certain m, then it is very easy to find an $L(M)$-formula for $d_p\varphi$: since $\varphi(x;b) \in p$ if and only if $m \vDash \varphi(x;b)$, we can set[14] $d_{p(x)}\varphi(x;y) = \varphi(m;y)$.

Observe that *the fact that p is realized in M is only needed in order for $\varphi(m,y)$ to be an $L(M)$-formula.* In fact, if $a \in \mathfrak{U}$ is a realization of p, then $b \in d_{p(x)}\varphi(x;y) \iff b \in M \wedge b \vDash \varphi(a;y)$, but the parameter a lives outside M. We are then led to explore the following concept:

Definition 1.49. Let M be a small model. A subset of M is *externally definable* iff it can be written as $M \cap \varphi(x;b) = \varphi(M;b) = \{a \in M \mid \mathfrak{U} \vDash \varphi(a;b)\}$ for some $\varphi(x;b) \in L(\mathfrak{U})$.

We will soon make the connection between defining schemata and externally definable subsets precise and show that the previous definition does not depend on $\mathfrak{U}$. Before that, let us see a couple of examples.

[13] The fact that this happens for all $M \vDash T$ is equivalent to the *stability* of T: see [Pil83, Corollary 1.21]. It can happen that all types over a certain model are definable but the theory is not stable; for instance this happens in $\mathbb{R} \vDash \mathsf{RCF}$, because of Theorem 1.52, o-minimality, and Dedekind completeness. See [vdD86] for the result and [MS94] for generalizations.

[14] Notice how variables are now "changing status", as we think of x as a parameter variable and y as an object variable.

Example 1.50. In $(\mathbb{Q}, <) \vDash \mathsf{DLO}$, the set $\{x \in \mathbb{Q} \mid x < \sqrt{2}\}$ is not definable, as can be easily shown by quantifier elimination, and this implies[15] that $\operatorname{tp}(\sqrt{2}/\mathbb{Q})$ is not $\mathbb{Q}$-definable. Anyway it is externally definable by the formula $x < \sqrt{2}$, or by the formula $x < b$ for any $b \in \mathfrak{U}$ with $b \equiv_{\mathbb{Q}} \sqrt{2}$.

One could wonder whether *all* subsets are externally definable. As the presence of Definition 1.49 may hint, the answer is in general *no*:

Example 1.51. By quantifier elimination, $\mathbb{Z}$ is not even externally definable in $(\mathbb{Q}, <)$.

Theorem 1.52. Let $A \subseteq M$. Then all externally definable subsets of M are A-definable if and only if all $p \in S(M)$ are A-definable.

Proof. The main point is the trivial observation that $\{b \in M \mid \varphi(x; b) \in \operatorname{tp}(a/M)\}$ is defined by $\varphi(a; y)$, and it is exactly $d_{\operatorname{tp}(a/M)}\varphi$. We feel free to spell out some details.

($\Rightarrow$) Let $p \in S_x(M)$ and $\varphi(x; y) \in L$; we want to find an $L(A)$ formula $d_p\varphi$ defining the set $X = \{b \in M \mid \varphi(x; b) \in p\}$. Fix a realization $a \vDash p$. Then obviously $X = \varphi(a; M)$. By hypothesis there is an $L(A)$-formula $\psi(y)$ such that $\varphi(a; M) = X = \psi(M)$, and setting $d_p\varphi = \psi$ proves that p is A-definable.

($\Leftarrow$) Let $X = \varphi(a; y)$, for some $a \in \mathfrak{U}$ and $\varphi(x; y) \in L$. If $p = \operatorname{tp}(a/M)$, by hypothesis there is a formula $d_p\varphi \in L(A)$ such that for all $b \in M$ we have $\varphi(x; b) \in p$ if and only if $\vDash (d_p\varphi)(b)$. But this means exactly $\varphi(a; M) = (d_p\varphi)(M)$. $\square$

Notation 1.53. $A \subset^+ N$ means that $A \subseteq N$ and N is $|A|^+$-saturated and $|A|^+$-strongly homogeneous[16]. If $A = M$ is a model we also write $M \prec^+ N$ with the same[17] meaning.

Definition 1.54. We say that $N \succ M$ *codes all externally definable subsets of* M iff for each externally definable $X \subseteq M$ there is $\psi(x)$ in $L(N)$ such that $\psi(M) = X$.

Proposition 1.55. If $M \prec^+ N$ then N already codes all externally definable subsets of M, i.e. every externally definable subset of M has an external definition with parameters in N.

Proof. If $\varphi(x; b)$ externally defines a subset of M and $b \equiv_M c$, then $\varphi(M; b) = \varphi(M; c)$ by definition. By saturation, such a c can be found in N. $\square$

[15] Or "is due to the fact", depending on viewpoints. See Theorem 1.52.

[16] Often we will only need saturation, but for the sake of precision it is better to fix notations once and for all.

[17] Recall that, by monster conventions as fixed in Section A.6, all inclusions between models are automatically elementary embeddings.

Corollary 1.56. Whether a subset is externally definable or not does not depend on $\mathfrak{U}$.

We could then be tempted to "make externally definable sets definable by fiat", i.e. to use the following construction:

Definition 1.57. Suppose that $N \succ M$ codes all externally definable subsets of M, and let $L^{\mathrm{ext}} = L \cup \{R_{\varphi(x;b)}(x) \mid \varphi(x;b) \in L(N)\}$. *Shelah's expansion* M^{ext} is the expansion of M to an L^{ext}-structure obtained by interpreting each $R_{\varphi(x;b)}(x)$ as $\varphi(M;b)$.

Some authors use the notation M^{Sh} instead. The model N is not explicitly mentioned in the notation M^{ext} for the following reason: although, strictly speaking, even the language L^{ext} depends on N, if $N_0, N_1 \succ M$ both code all externally definable subsets of M and L_0^{ext}, L_1^{ext} are the corresponding languages, then M_0^{ext} is the same structure of M_1^{ext} up to replacing every $R_{\varphi(x;b)}(x) \in L_0^{\mathrm{ext}}$ with $R_{\varphi(x;c)}(x) \in L_1^{\mathrm{ext}}$ for a suitable $\varphi(x;c) \in L(N_1)$. Thus, by Proposition 1.55, one may simply assume N to be the monster $\mathfrak{U}$; in any case, for most calculations we will need to fix such an N.

The fact that, if $\varphi \in L(M)$, then $d_p\varphi \in L^{\mathrm{ext}}$, suggests that in order to have a semigroup structure we may better work with $S_G(M^{\mathrm{ext}})$ instead of $S_G(M)$; in practice there is still another issue to be addressed, and it is precisely the same issue that we wanted to solve by passing to $S_G(M^{\mathrm{ext}})$, i.e. we want $d_p\varphi \in L^{\mathrm{ext}}$ *even for p an L^{ext}-type and φ an L^{ext}-formula.* In other words, now that the language has been expanded, all the L-externally definable subsets have become definable, but new L^{ext}-externally definable subsets may arise: there may be $L^{\mathrm{ext}^{\mathrm{ext}}}$-definable sets that are not L^{ext}-definable. Let us see why.

We could try to show, by induction on formulas, that every externally definable subset of M^{ext} is already definable. Nor conjunction, since $R_{\varphi\wedge\psi} \equiv R_\varphi \wedge R_\psi$, nor negation, since $R_{\neg\varphi} \equiv \neg R_\varphi$, present problems and, most important, $R_{\varphi(x;y)}(x;b) = R_{\varphi(x;b)}(x)$. A difficulty shows up when dealing with quantifiers. For instance, if $n \in N \setminus M$, then $M^{\mathrm{ext}} \not\models \exists x\ R_{x=n}(x)$ but $M \models R_{\exists x\ x=n}$, so a naïve approach will not work. Let us take a closer look.

Unraveling definitions we discover that, if $b \in M$, and N is used to define L^{ext} we have, inside M^{ext},

$$b \models \exists x\ R_{\varphi(x;y)}(x;y) \Leftrightarrow \exists m \in M\ b \models R_{\varphi(x;y)}(m;y) \Leftrightarrow \exists m \in M\ N \models \varphi(m;b)$$

$$b \models R_{\exists x\ \varphi(x;y)}(y) \Longleftrightarrow N \models \exists x\ \varphi(x;b) \Longleftrightarrow \exists n \in N\ N \models \varphi(n;b)$$

Thus $\exists x\ R_{\varphi(x;y)}(x;y)$ and $R_{\exists x\ \varphi(x;y)}(y)$ may define different subsets of M, the first one always being included in the second one. If the reverse inclusion always holds, then M^{ext} eliminates quantifiers, and our plan to use $S_G(M^{\mathrm{ext}})$ instead of $S_G(M)$ can be carried out. Unfortunately, this is not always the case. We will see in Theorem 3.56 that nonetheless, if we are working in a

NIP theory, then $\mathrm{Th}(M^{\mathrm{ext}})$ will indeed have quantifier elimination. For the moment we do not assume NIP (also because we did not even define what it means yet) and employ the solution adopted in [New12].

External Types

Since in dealing with $\mathrm{Def}_G(M^{\mathrm{ext}})$ we had problems with quantifiers, we turn our attention to another Boolean algebra.

Definition 1.58. $\mathrm{Def}_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ is the Boolean algebra of subsets of M^{ext} which are definable with a quantifier-free formula. Its Stone space will be denoted with $S_G^{\mathrm{qf}}(M^{\mathrm{ext}})$. The Boolean algebra of externally definable subsets of M will be denoted by $\mathrm{Def}_G^{\mathrm{ext}}(M)$, and its Stone space with $S_G^{\mathrm{ext}}(M)$. We put on both Stone spaces a $G(M)$-flow structure in the following way

$$g \cdot p(x) \ni \varphi(x) \iff p(x) \ni g^{-1} \cdot \varphi(x)$$

Lemma 1.59. $\mathrm{Def}_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ is isomorphic to $\mathrm{Def}_G^{\mathrm{ext}}(M)$, and $S_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ and $S_G^{\mathrm{ext}}(M)$ are isomorphic $G(M)$-flows.

Proof. The isomorphism of Boolean algebras and the homeomorphism between the corresponding Stone spaces is simply due to the fact that R_{-} commutes with connectives; the $G(M)$-flow structure is trivially preserved. □

Permanent Assumption 1.60. From now on we will use Lemma 1.59 tacitly and identify externally definable subsets of M and quantifier-free definable subsets of M^{ext}.

The importance of the following lemma should, at this point, be very clear.

Lemma 1.61 ([New12, Lemma 1.2])**.** If $p \in S_G^{\mathrm{ext}}(M)$ and $\varphi(x; b) \in \mathrm{Def}_G^{\mathrm{ext}}(M)$, then $\delta_p\varphi \in \mathrm{Def}_G^{\mathrm{ext}}(M)$.

Proof. Let $a \vDash p$ in $N_0 \ {}^{+}\!\succ N$, where N is used to code externally definable subsets of M. For all $g \in G(M)$ we have

$$\begin{aligned} g \vDash \delta_{p(x)}\varphi(x; b) &\iff \varphi(g \cdot x; b) \in p \\ &\iff N_0 \vDash \varphi(g \cdot a; b) \iff g \in \varphi(z \cdot a; b) \cap G(M) \end{aligned}$$

This shows that $\delta_p\varphi$ is externally definable[18]. □

Remark 1.62. The careful reader may have noticed that $\varphi(g \cdot x)$ is *not* a quantifier-free formula, even if $\varphi(x)$ is, because it is a shorthand for $\exists z\ \psi(g, x, z) \wedge \varphi(z)$, where ψ defines the group product. So, if she is thinking in terms of quantifier-free M^{ext}-types, she may wonder what it means

[18] But notice that a may be outside N, so we may have to replace $\varphi(z \cdot a; b)$ with a different formula if we want an external definition in $L(N)$.

for $\varphi(g \cdot x)$ to be an element of p. We are going now to explain this abuse of notation. If x ranges in $G(M)$ such a z as above always exists (and is unique) in $G(M)$. Hence, we have

$$M^{\text{ext}} \vDash \forall x \big(\exists z\, R_{\psi(g,x,z)\wedge\varphi(z)}(x,z) \leftrightarrow R_{\exists z\ \psi(g,x,z)\wedge\varphi(z)}(x)\big)$$

Since $R_{\exists z\ \psi(g,x,z)\wedge\varphi(z)}(x)$ *is* a quantifier free formula we can substitute it in place of $\varphi(g \cdot x)$ and meaningfully ask whether it belongs to $p \in S_G^{\text{qf}}(M)$.

Proposition 1.63 ([New12, Lemma 1.3])**.** δ_p commutes with Boolean operations and left multiplication by elements of $G(M)$.

Proof. It is very easy to check the first statement. As for the second one,

$$g \vDash \delta_p(h\varphi) \Leftrightarrow h\varphi \in gp \Leftrightarrow \varphi \in h^{-1}gp \Leftrightarrow h^{-1}g \vDash \delta_p\varphi \Leftrightarrow g \vDash h\delta_p\varphi \qquad \square$$

We are now ready to formally define our semigroup operation on $S_G^{\text{qf}}(M^{\text{ext}})$.

Definition 1.64. Let $p, q \in S_G^{\text{qf}}(M^{\text{ext}})$. We define $p \cdot q \in S_G^{\text{qf}}(M^{\text{ext}})$ as

$$\varphi \in p \cdot q \iff \delta_q\varphi \in p$$

Lemma 1.65 ([New12, Lemma 1.4])**.** $\delta_{q\cdot r} = \delta_q \circ \delta_r$. Moreover the operation $\cdot$ is associative.

Proof. Fix $g \in G(M)$ and, using Proposition 1.63, calculate

$$g \vDash \delta_{q\cdot r}\varphi \Leftrightarrow g^{-1}\varphi \in q \cdot r \Leftrightarrow \delta_r(g^{-1}\varphi) \in q \Leftrightarrow g^{-1}\delta_r\varphi \in q \Leftrightarrow g \vDash \delta_q(\delta_r\varphi)$$

and the "moreover" part follows:

$$\varphi \in (p \cdot q) \cdot r \Leftrightarrow \delta_r\varphi \in p \cdot q \Leftrightarrow \delta_q(\delta_r\varphi) \in p \Leftrightarrow \delta_{q\cdot r}\varphi \in p \Leftrightarrow \varphi \in p \cdot (q \cdot r) \quad \square$$

We abuse the notation a little more to make the parallelism with the previous section apparent.

Proposition 1.66 ([New12, Lemma 1.5])**.** Let $\vartheta_p(q) = p \cdot q$. Then

1. $\vartheta_p = \lim_{g\to p} \vartheta_g$
2. $\vartheta_- \colon S_G^{\text{qf}}(M^{\text{ext}}) \to E(S_G^{\text{qf}}(M^{\text{ext}}))$ is an isomorphism of semigroups and of G-flows.
3. For all $q \in S_G^{\text{qf}}(M^{\text{ext}})$ the map $- \cdot q$ is continuous.

Proof.

1. Notice that p is not an ultrafilter on the whole $\mathscr{P}(G(M))$; anyway our limit exists because

   $$\vartheta_p(q) \in [\varphi] \Leftrightarrow \varphi \in \vartheta_p(q) \Leftrightarrow \delta_q\varphi \in p \Leftrightarrow \{g \in G(M) \mid \varphi \in gq = \vartheta_g(q)\} \in p$$

   Since this happens for all φ and q, if we think of p as a filter on $\mathscr{P}(G(M))$ this says precisely that $\lim_{g\to p} \vartheta_g = \vartheta_p$.

2. The only thing that prevents us from recycling the proofs of Theorem 1.39 and Corollary 1.40 is the fact that it uses Lemma 1.38, which in turn depends on Lemma 1.16. In other words, we only need to show that ϑ_- is surjective, i.e. that for all $\mathcal{U} \in \beta G$, there is some $p \in S_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ such that $\lim_{g\to\mathcal{U}} \vartheta_g = \lim_{g\to p} \vartheta_g$. Unpackaging definitions we find that taking $p = \mathcal{U} \restriction \mathrm{Def}_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ works, the key point being the fact that all the $d_q\varphi$ are still externally definable by Lemma 1.61.

3. We have to check that $(-\cdot q)^{-1}[\varphi]$ is open. This is the set of the p such that $\varphi \in p \cdot q$, that coincides with $[d_q\varphi]$ by definition, which is open again by Lemma 1.61. □

We have therefore shown that $S_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ is isomorphic to its own envelope, and described the image of the composition under this isomorphism, which endows $S_G^{\mathrm{qf}}(M^{\mathrm{ext}})$ with and Ellis semigroup structure.

Heirs and Coheirs

Since in dealing with types it is often useful to realize them, we want to describe the realizations of $p \cdot q$ in terms of the ones of p and q. In order to do this we need some more model-theoretic tools.

Definition 1.67. Let $A \subseteq B$. A type p with parameters from B is *finitely satisfiable in* A iff for all $\varphi(x) \in p$ there is $a \in A$ such that $\vDash \varphi(a)$. The space of all such types in variables x is denoted $S_x^{\mathrm{fs}}(B, A)$.

Proposition 1.68. $S_x^{\mathrm{fs}}(B, A)$ is a closed subspace of $S_x(B)$.

Proof. A type is finitely satisfiable in A if and only if it omits all formulas that have no point in A, i.e. if and only if it belongs to the closed subset $\bigcap_{\varphi(A)=\emptyset}[\neg\varphi(x)]$. □

Of course if $A = B = M$ is a model then $S_x^{\mathrm{fs}}(M, M) = S_x(M)$. It is *not* true that all types over A are finitely satisfiable in A: take $A = \emptyset$, for example. We now give a name to finitely satisfiable extensions.

Definition 1.69. Let $A \subseteq B$ and $p \in S_x(A)$. A *coheir of p over B* is any $q \in S_x^{\mathrm{fs}}(B, A)$ such that $q \restriction A = p$. In other words, it is an extension of p to a type with parameters in B that is finitely satisfiable in A.

Proposition 1.70. If $p \in S_x^{\mathrm{fs}}(B, A)$ and $C \supseteq B$, then p can be extended to an element of $S_x^{\mathrm{fs}}(C, A)$. In particular this always holds when $B = A = M$ is a model.

Proof. Since for $\varphi \in p$ the set $\varphi(A)$ is never empty, we can extend $\{\varphi(A) \mid \varphi \in p\}$ to an ultrafilter $\mathcal{U} \in \beta A$. Then, given any $C \supseteq B$, define a type $q_{\mathcal{U}} \in S_x^{\mathrm{fs}}(C, A)$ by $\varphi(x; b) \in q_{\mathcal{U}} \iff \varphi(A; b) \in \mathcal{U}$. □

As the "co" in "coheir" may hint, there is a notion called *heir*, and the two are — in a sense — dual.

Definition 1.71. If $A \subseteq B$, $p \in S_x(A)$ and $q \in S_x(B)$, we say that q is an *heir* of p iff for all $\varphi(x; y) \in L$ and $b \in B$, if $\varphi(x; b) \in q$ there is $a \in A$ such that $\varphi(x; a) \in p$.

Proposition 1.72 (Heir-coheir duality)**.** $\mathrm{tp}(a/Ab)$ is an heir of $\mathrm{tp}(a/A)$ if and only if $\mathrm{tp}(b/Aa)$ is a coheir of $\mathrm{tp}(b/A)$

Proof. ($\Rightarrow$) Let $\varphi(x; y) \in L(A)$ and suppose that $\vDash \varphi(a; b)$. By the "heir" hypothesis there is $c \in A$ such that $\vDash \varphi(a; c)$, and this shows that $\varphi(a; y)$ is finitely satisfiable in A.

($\Leftarrow$) Dually, let $\varphi(x; y) \in L(A)$ and suppose that $\vDash \varphi(a; b)$. By the "coheir" hypothesis $\varphi(a; y)$ is finitely satisfiable in A, and this shows that there is $c \in A$ such that $\vDash \varphi(a; c)$. □

Example 1.73. Let p be the type at $+\infty$ over some $M \vDash \mathsf{RCF}$ and let $a \vDash p$. It can be easily checked that if b realizes an heir of p over Ma, then "$b \gg a$", meaning that b is bigger than any point in the definable closure of a, i.e. of any polynomial in a, and if b realises a coheir of p over Ma then $b \ll a$. Notice that if we replace RCF with an o-minimal expansion of hers, then the definable closure of a will be bigger; for instance, if we consider $\mathrm{Th}(\mathbb{R}_{\mathrm{exp}})$ (which is o-minimal by [Wil96]) b will be bigger than any exponential polynomial involving a. In this special case "being an heir" captures the notion of being "more infinite". Similarly, if we repeat the construction with $p = \mathrm{tp}(0^+/M)$, it will capture the notion of being "more infinitesimal".

The following fact will not be used, but it is sufficiently standard to be worth mentioning.

Fact 1.74 ([Pil83, Proposition 1.17])**.** A type in $S_x(M)$ is definable if and only if for all $A \supseteq M$ it has exactly one heir in $S_x(A)$.

For the moment we put heirs aside, but we will need them in later chapters. Back to coheirs, we now explain why speaking of coheirs over a sufficiently big model is essentially the same as speaking of external types.

Theorem 1.75. Suppose[19] that N codes all externally definable subsets of M. Then every element of $S^{\mathrm{fs}}_x(N, M)$ has a unique M-finitely satisfiable global extension. In other words, the restriction $S^{\mathrm{fs}}_x(\mathfrak{U}, M) \to S^{\mathrm{fs}}_x(N, M)$ is injective (hence an homeomorphism).

Proof. One extension exists by Proposition 1.70. Let $p \neq q \in S^{\mathrm{fs}}_x(\mathfrak{U}, M)$, as witnessed by $\varphi \in p$, $\neg\varphi \in q$. By hypothesis both φ and $\neg\varphi$ have points in M, so they correspond to X, $X^{\complement}$, both non-empty and externally definable by some $\psi, \neg\psi \in L(N)$ by hypothesis. If we prove that $\psi \in p$ and similarly $\neg\psi \in q$ we are done, so suppose $\neg\psi \in p$. Then $\neg\psi \wedge \varphi$ is in p and so should have a point in M, but such a point would inhabit $X^{\complement} \cap X$. □

Corollary 1.76. If all externally definable subsets of M are definable in M (e.g. we are working in an M^{ext} that eliminates quantifiers, or in $\mathbb{R}$), then types over M have unique coheirs over arbitrary sets of parameters.

Proof. Even if we want to extend a type to a parameter set that lives outside $\mathfrak{U}$, it suffices to apply Theorem 1.75 to a suitable bigger monster $\widetilde{\mathfrak{U}}$ and then take a restriction. □

Proposition 1.77. Suppose that N codes all externally definable subsets of M. Then $S^{\mathrm{ext}}_x(M)$ is homeomorphic to $S^{\mathrm{fs}}_x(N, M)$.

Proof. It is enough to send $p \in S^{\mathrm{fs}}_x(N, M)$ to $\{\varphi(M) \mid \varphi \in p\}$. In the other direction, send $q \in S^{\mathrm{ext}}_x(M)$ to $\{\varphi(x) \in L(N) \mid \varphi(M) \in q\}$. □

Remark 1.78. If we define a $G(M)$-flow structure on $S^{\mathrm{fs}}_G(N, M)$ in the obvious way, i.e. $\varphi \in gp \iff g^{-1}\varphi \in p$, then the previous homeomorphism is also an isomorphism of G-flow. Hence we will freely identify $S^{\mathrm{ext}}_G(M)$ with $S^{\mathrm{fs}}_G(N, M)$.

We can now understand realizations of $p \cdot q$.

Theorem 1.79 ([New09, Lemma 4.1 (1)[20]])**.** Suppose that N codes all externally definable subsets of M and work in $S^{\mathrm{fs}}_G(N, M)$. Then $p \cdot q = \mathrm{tp}(a \cdot b/N)$ where b realizes q and a realizes the unique M-finitely satisfiable extension of p to Nb.

Proof. Denote said unique extension with $p \mid^{\mathrm{ch}}_{M} Nb$. Given $\varphi(x) \in p \cdot q$ we want to show that $\models \varphi(a \cdot b)$. Let $\psi(x) \in L(N)$ be any formula such that $\psi(x) \cap M = \delta_q \varphi$.

Claim. $\big(\psi(x) \leftrightarrow \varphi(x \cdot b)\big) \in p \mid^{\mathrm{ch}}_{M} Nb$.

[19] For example $N \mathrel{{}^{+}{\succ}} M$, by Proposition 1.55, or we could "be lucky", have all externals already M-definable and take $N = M$.

[20] Mutatis mutandis.

Proof of the Claim. If this is not the case, then by M-finite satisfiability the equivalence above fails for at least one $g \in G(M)$. But for such a g we have

$$g \vDash \psi(x) \iff g \in \delta_q \varphi \iff \varphi(x) \in gq \overset{\dagger}{\iff} \underbrace{\varphi(g \cdot x)}_{=g^{-1}\varphi(x)} \in q \iff \vDash \varphi(g \cdot b)$$

where the equivalence † is proven via the isomorphism with $S_G^{\text{ext}}(M)$ exploiting the fact that, being an isomorphism of $G(M)$-flows, it commutes with multiplication by $g \in G(M)$. □ CLAIM

Since $p \mid_M^{\text{ch}} Nb$ is an extension of p and $\psi \in p$ by definition of $p \cdot q$, it then follows from the Claim that $\varphi(x \cdot b) \in p \mid_M^{\text{ch}} Nb$, and this is equivalent to $a \vDash \varphi(x \cdot b)$. □

A natural question would now be to try to understand what happens if we change our model, i.e. to compare $S_G^{\text{ext}}(M)$ and $S_G^{\text{ext}}(N)$ and the relative Ellis groups when $M \preceq N$. This task is quite hard, and some results can be found in [New12]. Anyway, we will see as a corollary of the main result studied in this thesis that under further hypotheses the Ellis group does not depend on the model.

The Conjecture

We now have two groups attached to a couple (G, M): one is the interpretation $G(M)$ of G in M, the second is the Ellis group associated to the $G(M)$-flow $S_G^{\text{ext}}(M)$. There is a third construction that has been extensively studied and goes under the name of G/G^{00}.

Definition 1.80. Let κ be the saturation of $\mathfrak{U}$ and $|A| < \kappa$. The *A-type-connected component* G_A^{00} is defined as the intersection of all subgroups of $G(\mathfrak{U})$ which are type-definable with parameters in A and of bounded index, i.e. such that $[G(\mathfrak{U}) : G_A^{00}(\mathfrak{U})] < \kappa$.

Definition 1.81. We say that G^{00} *exists* iff for all small A it happens that $G_A^{00} = G_\emptyset^{00}$, and in this case we define G^{00} to be $G_\emptyset^{00}$.

We will later see (Proposition 2.54) that if M is a small model the projection to the quotient $G(\mathfrak{U}) \to G(\mathfrak{U})/G_M^{00}(\mathfrak{U})$ factors through a surjective $\pi \colon S_G(M) \to G/G_M^{00}$, and that G_M^{00} is normal (Proposition 2.56). This implies the following.

Remark 1.82. G/G_M^{00} does not depend on $\mathfrak{U}$, since $S_G(M)$ does not: to know what coset of $G_M^{00}(\mathfrak{U})$ one $g \in G(\mathfrak{U})$ belongs to it is sufficient to know its type over M, and using saturation it is easy to define an isomorphism between $G/G_M^{00}(\mathfrak{U}_0)$ and $G/G_M^{00}(\mathfrak{U}_1)$.

This enables us to compare the Ellis group associated to (G, M) with G/G^{00}_M, for the following reason. Since, by Remark 1.78, we can identify $S^{\mathrm{ext}}_G(M)$ with $S^{\mathrm{fs}}_G(\mathfrak{U}, M) \subseteq S_G(\mathfrak{U}, M)$, it makes sense to fix an Ellis group uI inside it and consider the restriction $\pi \restriction uI$.

Theorem 1.83 ([New09, Proposition 4.4], [CPS14, Remark 5.4])**.** The restriction $\pi\colon uI \to G/G^{00}_M$ is a surjective group homomorphism.

Proof. By Theorem 1.79, if $p, q \in S^{\mathrm{fs}}_G(\mathfrak{U}, M)$, $b \vDash q$ and $a \vDash p \mid^{\mathrm{ch}}_M \mathfrak{U}b$,

$$\pi(p \cdot q) = \pi(\mathrm{tp}(a \cdot b)/\mathfrak{U}) = (a \cdot b)G^{00}_M = (aG^{00}_M) \cdot (bG^{00}_M) = \pi(p) \cdot \pi(q)$$

so $\pi\colon S^{\mathrm{fs}}_G(\mathfrak{U}, M) \to G/G^{00}_M$ is an homomorphism of semigroups. In particular its restriction to uI is an homomorphism of groups. To show it is surjective, fix $gG^{00}_M \in G/G^{00}_M$ and let $p = \mathrm{tp}(g/M)$. Since I is an ideal and $u \in I$ we have that $pu \in I$, and so $upu \in uI$. Since π is an homomorphism and u is the identity of uI we have that $\pi(upu) = \pi(u)\pi(p)\pi(u) = \pi(p) = gG^{00}_M$. $\square$

A question arises: is $\pi \restriction uI$ an isomorphism? The answer is *no*, as a counterexample, namely $\mathrm{SL}(2, \mathbb{R})$, was found in [GPP15]. One could then try to impose restrictions on T, but the fact that there is already a counterexample in an o-minimal theory suggests that this may not be dependent solely on the tameness of T.

We can now state the conjecture, formulated by Pillay in [CPS14], around which this thesis is built.

Conjecture 1.84 (Ellis Group Conjecture)**.** If G is a definably amenable group and T is NIP, then $\pi \restriction uI$ is an isomorphism.

We have yet to explain what a NIP theory is and what it means for a group to be definably amenable. After some preliminary model-theoretic study in Chapter 2, we will define NIP theories in Chapter 3, and verify that in this context all the relevant concepts are preserved when passing to M^{ext}. In Chapter 4 we will turn our attention to definably amenable groups and examine how, after special cases were shown to hold in [New12,Pil13,CPS14], Chernikov and Simon proved in [CSed] that the Ellis Group Conjecture is true.

Topological dynamics relies vitally on notions such as the ones of continuous map, compactification, flow. In [GPP14] definable analogues of such notions were introduced, in order to develop a theory of "tame topological dynamics". See also [Pil13] for further results.

Chapter 2

Model-Theoretic Tools

In our study of NIP theories and definably amenable groups we will often use invariant types, forking, and Keisler measures. Under the NIP hypothesis measures will behave significantly better, while forking will have strong characterizations in terms of invariance; nonetheless they can be legitimately considered (and used!) without assuming it. For this reason, we will investigate these concepts before even defining NIP. This will have the additional advantage of making clear what facts depend on the NIP and what do not. References for this chapter are [Sim15, TZ12].

2.1 Invariant Types

As we saw in the previous chapter, if types have a "privileged" extension to any bigger set of parameters, this can be used to define operations on them in terms of their realizations. For instance, Theorem 1.79 is due to the ability of extending p in a way that "takes into account" a particular realization b of q. In that case, "privileged" meant "finitely satisfiable in M" which is another way of saying "a coheir of it restriction to M". As the reader may expect, one could employ heirs in a similar fashion, and since we want a "canonical" extension, Fact 1.74 suggests to start with definable types. In fact, this is probably one of their most essential features (cf. [Las76]).

Definition 2.1. Let $p \in S_x(M)$ be definable, and let $B \supseteq M$. Define $p \mid B \in S_x(B)$ in the following way: if $\varphi(x;y) \in L$ and $b \in B$, then $\varphi(x;b) \in p \mid B$ if and only if $b \models d_{p(x)}\varphi(x;y)$

In other words, since for all $a \in M$ whether $\varphi(x;a) \in p$ or not depends only on whether $a \models d_p\varphi$ holds, we define an extension of p to B imposing the same "rule" to hold for all $b \in B$. Then we can try to generalize this construction from definable types to types for which a reasonable "rule" exists. In other words, we could try to relax the hypotheses on $d_p\varphi$; the main point of *invariant types* is to allow $d_p\varphi$ to be any subset of the type space. Let us begin giving explicit definitions.

Generalizing Definable Types

Definition 2.2. Let $p \in S_x(B)$ and let $A \subseteq B$. We say that p is *A-invariant* iff, for all $\varphi(x; y) \in L$, if $b, c \in B$ and $b \equiv_A c$, then $\varphi(x; b) \leftrightarrow \varphi(x; c) \in p$. The subspace of A-invariant types in $S_x(B)$ is denoted $S_x^{\mathrm{inv}}(B, A)$. If $p \in S_x(\mathfrak{U})$ we say that p is *invariant* iff there is some small A such that p is A-invariant. We denote the space of such types with $S_x^{\mathrm{inv}}(\mathfrak{U})$.

In other words, $p(x) \in S_x(B)$ is A-invariant iff for every $\varphi(x; y) \in L$ and $b \in B$ whether $\varphi(x; b) \in p$ or not depends only on $\mathrm{tp}(b/A)$, thus it makes sense to lift $\{b \in B \mid \varphi(x; b) \in p\}$ to a subset of $S_y(A)$. We will often use tacitly the following two lemmas:

Lemma 2.3. If $A \subseteq B \subseteq C$ and $p \in S_x(C)$ is A-invariant, then it is B-invariant.

Proof. If $b \equiv_B c$, a fortiori $b \equiv_A c$. □

Lemma 2.4. If $p, q \in S_x^{\mathrm{inv}}(\mathfrak{U})$, then there is a small model M such that $p, q \in S_x^{\mathrm{inv}}(\mathfrak{U}, M)$.

Proof. If p is A_p-invariant and q is A_q-invariant, by Lemma 2.3 if suffices to take $A = A_p \cup A_q$ and consider any small $M \supseteq A$. □

Proposition 2.5. $S_x^{\mathrm{inv}}(B, A)$ is closed in $S_x(B)$.

Proof. By definition, $S_x^{\mathrm{inv}}(B, A) = \bigcap_{\substack{b \equiv_A c \\ \varphi(x;y) \in L}} [\varphi(x; b) \leftrightarrow \varphi(x; c)]$. □

Lemma 2.6. If $A \subset^+ M$, then $p \in S_x^{\mathrm{inv}}(M, A)$ if and only if p is fixed by all elements of $\mathrm{Aut}(M/A)$.

Proof. By definition, $\varphi(x; f(b)) \in f(p)$ iff $\varphi(x; b) \in p$, and since $A \subset^+ M$ we have $b \equiv_A c \iff \exists f \in \mathrm{Aut}(M/A)\ f(b) = c$. □

Remark 2.7. If we want to speak of A-invariant *partial* types, some attention is needed. if one says that $\pi(x)$ is invariant under $\mathrm{Aut}(\mathfrak{U}/A)$, then it means that if $\varphi(x; b) \in \pi(x)$ then for *all* $a \equiv_A b$ in the monster $\varphi(x; a) \in \pi(x)$; saying that if $a \equiv_A b$ then $\pi(x)$ cannot contain both $\varphi(x; a)$ and $\neg\varphi(x; b)$ is clearly weaker, since $\pi(x)$, being partial, could contain neither $\varphi(x; a)$ nor $\neg\varphi(x; a)$. For instance, saying that $p \in S_x(M)$ is A-invariant means that p is fixed by elements of $\mathrm{Aut}(M/A)$, but if we regard p as a global partial type this is weaker than being fixed by elements of $\mathrm{Aut}(\mathfrak{U}/A)$. We can avoid these subtleties by writing $p(x) \vdash \varphi(x; a) \leftrightarrow \varphi(x; b)$ instead of $\varphi(x; a) \leftrightarrow \varphi(x; b) \in p$, because then $p(x) = f(p) \vdash \varphi(x; f(a)) \leftrightarrow \varphi(x; f(b))$.

Invariant types are a common generalization of definable and of finitely satisfiable ones:

Proposition 2.8. *A-finitely satisfiable types are A-invariant.*

Proof. Suppose that $\varphi(x;b) \triangle \varphi(x;c) \in p \in S_x^{\mathrm{fs}}(B;A)$. By finite satisfiability there is $a \in A$ such that $\vDash \varphi(a;b) \triangle \varphi(a;c)$, hence $b \not\equiv_A c$ □

Corollary 2.9. *Let $M \subseteq B$. Every type in $S_x(M)$ has an extension in $S_x^{\mathrm{inv}}(B, M)$.*

Proof. Every type has one in $S_x^{\mathrm{fs}}(B, M)$ by Proposition 1.70, and by Proposition 2.8 $S_x^{\mathrm{fs}}(B, M) \subseteq S_x^{\mathrm{inv}}(B, M)$. □

Anyway the inclusion may be strict:

Example 2.10. In DLO we have that $\mathrm{tp}(+\infty/\mathfrak{U})$ is $\emptyset$-definable, hence, by the following proposition, $\emptyset$-invariant, but it is not finitely satisfiable in any small set.

Proposition 2.11. *A-definable types are A-invariant.*

Proof. If $b \equiv_A c$, since every $d_p\varphi$ is in $L(A)$, we have $\vDash d_p\varphi(b) \leftrightarrow d_p\varphi(c)$. □

Hence $p(x)$ is A-definable iff not only it makes sense to lift $\{b \in B \mid \varphi(x;b) \in p\}$ to a subset of $S_y(A)$, but such a set is clopen, since can be written as $[d_p\varphi]$ because $d_p\varphi \in L(A)$.

Definition 2.12. In analogy to the case of definable types, if $p \in S_x(B)$ is A-invariant, we define

$$d_{p(x)}\varphi(x;y) = \{q \in S_y(A) \mid \exists b \in B \; b \vDash q \wedge \varphi(x;b) \in p\}$$

Remark 2.13. These objects behave best if[1], for all n, all types in $S_n(A)$ are realized in B: in this case $S_y(A) = d_p\varphi \cup d_p\neg\varphi$. Otherwise, if there is no $b \vDash q$ in B, the union $d_p\varphi \cup d_p\neg\varphi$ does not contain q.

The following proposition says that, under suitable hypotheses[2], Definition 2.1 can be generalized to invariant types.

Theorem 2.14 ([Sim15, discussion at page 19])**.** *Let $A \subseteq M$ and suppose that[3], for all n, all types in $S_n(A)$ are realized in M. Let $p \in S_x^{\mathrm{inv}}(M, A)$ and $C \supseteq M$. Then*

$$p \mid C = \{\varphi(x;b) \mid \varphi(x;y) \in L, b \in C, \mathrm{tp}(b/A) \in d_p\varphi\}$$

is the unique extension of p to an element of $S^{\mathrm{inv}}(C, A)$.

Proof. We have quite a number of things to check about $p \mid C$:

---

[1]. . . and some books only define them when. . .

[2]To know what happens if we try to remove them see Remark 2.16.

[3]For instance $A \subset^+ M$, e.g. $M = \mathfrak{U}$ and A is small.

Well-definedness Asking if $\mathrm{tp}(b/A) \in d_p\varphi$ makes sense by A-invariance.

Consistency In order to have both $\varphi(x; b)$ and $\neg\varphi(x; b)$ in $p \mid C$ it is necessary to have $b_0, b_1 \in M$, both realizing $\mathrm{tp}(b/A)$, such that $\varphi(x; b_0) \wedge \neg\varphi(x; b_1) \in p$. Since p is invariant this cannot happen.

Completeness By Remark 2.13 either $\mathrm{tp}(b/A) \in d_p\varphi$ or $\mathrm{tp}(b/A) \in d_p\neg\varphi$.

Invariance $p \mid C$ is A-invariant by definition.

Uniqueness Since we already proved that $p \mid C$ is a complete type, if p_0, p_1 are both candidates to be $p \mid C$ then there is $\varphi(x; b)$ such that $\varphi(x; b) \in p_0$ and $\neg\varphi(x; b) \in p_1$. But then $\mathrm{tp}(b/A) \in d_p\varphi \cap d_p\neg\varphi = \emptyset$. □

Corollary 2.15 (Cf. Theorem 1.75)**.** Under the same hypotheses, the restriction $S^{\mathrm{inv}}_x(B, A) \to S^{\mathrm{inv}}_x(M, A)$ is an homeomorphism.

Remark 2.16. If M omits some $q \in S_y(A)$, the construction of Theorem 2.14 only yields a partial type (cf. Remark 2.13). Hence in this case p may have more than one A-invariant extension.

Since we defined a global type to be invariant iff it is A-invariant for some small A, we will better check that this construction does not depend on A, in order to apply it to invariant types *tout-court*.

Proposition 2.17. If $p \in S_x(\mathfrak{U})$ is invariant over two small sets A_0 and A_1 and $C \supseteq \mathfrak{U}$, then $p \mid_{A_0} C = p \mid_{A_1} C$.

Proof. Fix any $c \in C$. If $b \in \mathfrak{U}$ is any realization of $\mathrm{tp}(c/A_0 \cup A_1)$ we have both $c \equiv_{A_0} b$ and $c \equiv_{A_1} b$. Hence

$$\varphi(x; c) \in p \mid_{A_0} C \iff \varphi(x; b) \in p \iff \varphi(x; c) \in p \mid_{A_1} C$$ □

The previous results allows us to think of global invariant types as the "truly global" ones, since as long as there is a defining schema $d_p \colon \{\varphi(x; y) \mid \varphi \in L\} \to S_y(M)$ of $p \in S^{\mathrm{inv}}_x(\mathfrak{U}, M)$, the monster $\mathfrak{U}$ does not even matter: even if we replace it with a larger $\widetilde{\mathfrak{U}} \mathrel{{}^{+}\!\succ} \mathfrak{U}$, the schema d_p will "carry" p to $S^{\mathrm{inv}}_x(\widetilde{\mathfrak{U}})$. Notice that by Lemma 2.3 we may also see $d_p\varphi$ as a subset of $S_x(\mathfrak{U})$.

Example 2.18. Let us see what an element of $S^{\mathrm{inv}}_1(\mathfrak{U})$ looks like in an o-minimal theory. Realized types are always invariant over a small model containing the realization. By o-minimality, non-realized types in one variable correspond to non-realized cuts, i.e. a type $p(x) \in S_1(\mathfrak{U}) \setminus \mathfrak{U}$ is determined by $A = \{a \in \mathfrak{U} \mid x > a \in p\}$ and $B = \{b \in \mathfrak{U} \mid x < b \in p\}$. Then p is invariant if and only if at least one between the cofinality of A and the coinitiality of B is small.

Proof. If one between A and B is empty, then p is a type at ∞, and so it is $\emptyset$-definable, hence $\emptyset$-invariant. Suppose that $A, B \neq \emptyset$ and there is a small increasing sequence $(a_i)_{i\in I}$ which is cofinal in A. If $M \supseteq (a_i)_{i\in I}$ is a small model, then p is M-invariant because it is finitely satisfiable in M. The symmetric argument works for B.

Now suppose that neither the cofinality of A nor the coinitiality of B are small, but for some small M we have $p \in S_1^{\mathrm{inv}}(\mathfrak{U}, M)$. Notice that $A < B$ and that since p is complete there can be no point between A and B. Then, by cofinality reasons, we can find $a \in A$ and $b \in B$ such that no point of M lies between a and b, and this implies $a \equiv_M b$ and contradicts M-invariance of p. □

Hence in this case if a type in dimension 1 fails to be invariant it is because the information it contains relies critically on monster-many parameters, and the monster "does not know" how to extend it canonically to larger models. In a sense, the information contained in an invariant type is small enough "to let $\mathfrak{U}$ communicate it".

Remark 2.19. If $M \prec^+ N$ and $p \in S_x^{\mathrm{fs}}(N, M)$ we have defined two "canonical extensions" of p to a type with parameters in $B \supseteq N$: the one from Theorem 1.75 and the one from Theorem 2.14. The two extensions are indeed the same: since $p \mid_M^{\mathrm{ch}} B$ is finitely satisfiable in M, it is in particular M-invariant. Thus, by Theorem 2.14, it coincides with $p \mid B$, since the latter is the *unique* M-invariant extension of p to B.

Corollary 2.20. If $p \in S^{\mathrm{inv}}(\mathfrak{U})$ is finitely satisfiable in a small M, the same is true for $p \mid C$.

Proof. $p \mid C = p \mid_M^{\mathrm{ch}} C$ and the latter is finitely satisfiable in M. Alternatively and withouth mentioning $\mid_M^{\mathrm{ch}}$, let $\varphi(x; c) \in p \mid C$ and find $b \in \mathfrak{U}$ such that $b \equiv_M c$ and $\varphi(x; b) \in p$. Then let $m \in M$ be such that $\vDash \varphi(m; b)$ and, since $b \equiv_M c$, we have $\vDash \varphi(m; c)$. □

Example 2.21. Take a small $M \vDash \mathsf{DLO}$ and let $p = \mathrm{tp}(M^+/\mathfrak{U}) \in S_x^{\mathrm{fs}}(\mathfrak{U}, M)$ be "the type of a point just to the right of M", i.e. the unique complete type extending

$$\{x > a \mid \exists m \in M\ a < m\} \cup \{x < b \mid \forall m \in M\ m < b\}$$

Let $c \vDash p$. Then, as $c > M$, $(x < c) \in p \mid \mathfrak{U}c$.

Lemma 2.22. If $C \supseteq \mathfrak{U}$ and $p \in S_x^{\mathrm{inv}}(\mathfrak{U})$, then $p \mid C$ is an heir of p.

Proof. Let M be a small model such that p is M-invariant and suppose that $\varphi(x; c) \in p \mid C$. By saturation there is $b \in \mathfrak{U}$ such that $b \equiv_M c$, and by definition $\varphi(x; b) \in p$. □

Products

Definition 2.23. Let $p \in S_x^{\mathrm{inv}}(\mathfrak{U})$ and $q \in S_y(\mathfrak{U})$. Let $b \models q$ and $a \models p \mid \mathfrak{U}b$. We define $p(x) \otimes q(y)$ as $\mathrm{tp}(a, b/\mathfrak{U})$

If we want to avoid realizing global types (because realizations will generally live outside the monster), we can use one of the following three characterizations of $\otimes$.

Lemma 2.24. Let $p \in S_x^{\mathrm{inv}}(\mathfrak{U}, A)$ and $q \in S_y(\mathfrak{U})$. Then $\varphi(x, y) \in p(x) \otimes q(y)$ if and only if $(q(y) \restriction A) \in d_p\varphi$.

Proof. Let $(a, b) \models p(x) \otimes q(y)$. By definition $\varphi(x, y) \in p(x) \otimes q(y)$ if and only if $\varphi(x, b) \in p \mid \mathfrak{U}b$, if and only if $\mathrm{tp}(b/A) \in d_p\varphi$. But $\mathrm{tp}(b/A) = q \restriction A$. □

Lemma 2.25. Suppose that $p \in S_x^{\mathrm{inv}}(\mathfrak{U}, A)$, $q \in S_y(\mathfrak{U})$, $B \supseteq A$ is a small set and $\varphi(x, y) \in L(B)$. Then $\varphi(x, y) \in p \otimes q$ if and only if for all $b \models q \restriction B$ living in $\mathfrak{U}$ we have $\varphi(x, b) \in p$.

Proof. Let c be the parameters from $B \setminus A$ appearing in $\varphi(x, y) = \varphi(x, y; c)$ and let $\tilde{b} \models q$. Since p is A-invariant, it is a fortiori B-invariant, so $\varphi(x, y; c) \in p \otimes q$ if and only if for all $b \equiv_B \tilde{b}$ living in $\mathfrak{U}$ we have $\varphi(x, b; c) \in p$. □

Corollary 2.26. Suppose that $p \in S_x^{\mathrm{inv}}(\mathfrak{U}, A)$, $q \in S_y(\mathfrak{U})$ and $B \supseteq A$ is a small set. Then $(a, b) \models p \otimes q \restriction B$ if and only if $b \models q \restriction B$ and $a \models p \restriction Bb$.

Proof. This is just a restatement of Lemma 2.25: if $(a, b) \models p \otimes q \restriction B$ then $b \models q \restriction B$ and, if $\varphi(x, y)$ is an $L(B)$ formula, then by the lemma $a \models \varphi(x, b) \iff \varphi(x, y) \in p \otimes q \iff \varphi(x, b) \in p \restriction Bb$. Conversely, suppose $b \models q$ and $a \models p \restriction Bb$. If $\varphi(x, y) \in p \otimes q$, then by the lemma $\varphi(x, b) \in p$, and so $a \models \varphi(x, b)$. □

Lemma 2.27. If both $p(x)$ and $q(y)$ are invariant, then $p(x) \otimes q(y)$ is too.

Proof. Let p and q be M-invariant, and fix $\widetilde{\mathfrak{U}} \ {}^{+}\!\succ \mathfrak{U}$ such that $\widetilde{\mathfrak{U}} \ni (a, b) \models p \otimes q$. Let $f \in \mathrm{Aut}(\mathfrak{U}/M)$ and extend it to $\mathrm{Aut}(\widetilde{\mathfrak{U}}/M)$ by strong homogeneity. We have $f(p \otimes q) = f(\mathrm{tp}(a, b/\mathfrak{U})) = \mathrm{tp}((f(a), f(b))/\mathfrak{U})$, so if we show that $(f(a), f(b)) \models p \otimes q$ we are done. Fix $\varphi(x, y) \in p \otimes q$. Since q is M-invariant, we have $\mathrm{tp}(f(b)/\mathfrak{U}) = f(q) = q \in d_p\varphi$. Hence, $a \models p \mid \mathfrak{U}f(b)$. Since $p \mid \mathfrak{U}f(b)$ is still M-invariant, $f(a) \models f(p \mid \mathfrak{U}f(b)) = p \mid \mathfrak{U}f(b)$. This shows that $(f(a), f(b)) \models p \otimes f(q) = p \otimes q$. □

Proposition 2.28. $\otimes$ is associative.

Proof. Suppose that $p(x)$, $q(y)$ and $r(z)$ are A-invariant global types, and let $c \models r$, $b \models q \mid \mathfrak{U}c$, and $a \models p \mid \mathfrak{U}bc$. Then $\mathrm{tp}(a, b, c/\mathfrak{U}) = p(x) \otimes (q(y) \otimes r(z))$ by definition, so it suffices to check that $(a, b) \models p(x) \otimes q(y) \mid \mathfrak{U}c$. This is easy: $\varphi(x, y, c) \in p(x) \otimes q(y) \mid \mathfrak{U}c$ if and only if there is some $d \equiv_A c$ in $\mathfrak{U}$ such that

$\varphi(x, y, d) \in p \otimes q$. Since $b \models q$, this happens if and only if $\varphi(x, b, d) \in p \mid \mathfrak{U}b$, if and only if $\models \varphi(a, b, d)$ because $a \models p \mid \mathfrak{U}bc \supseteq p \mid \mathfrak{U}b$. By the previous lemma $\operatorname{tp}(a, b, c/\mathfrak{U})$ is A-invariant, so $\operatorname{tp}(a, b/\mathfrak{U})$ is too[4]. But then, since $d \equiv_A c$, we have $\varphi(x, y, c) \in \operatorname{tp}(a, b) \mid \mathfrak{U}c$. □

Lemma 2.29. The product of M-finitely satisfiable global types is M-satisfiable.

Proof. If $\varphi(x; y) \in p(x) \otimes q(x)$, let $b \models q$. Then $\varphi(x; b) \in p \mid \mathfrak{U}b$ by definition. Since p is finitely satisfiable in M, the same is true of $p \mid \mathfrak{U}b$ by Corollary 2.20. But if $m \in M$ is such that $\models \varphi(m, b)$, then $\varphi(m, y) \in \operatorname{tp}(b/\mathfrak{U}) = q$, which is finitely satisfiable in M, hence we find $n \in M$ such that $\models \varphi(m, n)$. □

Remark 2.30. $\otimes$ on (invariant) types is analogous to $\otimes$ on ultrafilters (cf. Definition 1.19), except it is written *backwards*: compare with Theorem 1.79, where the product $\cdot$ was defined in a way "compatible with ultrafilters", i.e. backwards with respect to the product on invariant types.

Remark 2.31. It is not true that if p and q are both invariant then $(a, b) \models p(x) \otimes q(y)$ if and only if $a \models p(x)$ and $b \models q \mid \mathfrak{U}a$; in other words, Theorem 1.79 does not generalize. This is due to the fact that in general $p(x) \otimes q(y) \neq q(y) \otimes p(x)$.

An attempt to adapt the proof of Theorem 1.79 in a straightforward way will fail because $d_q\varphi$[5] need not be clopen, i.e. q need not be definable: in other words there is no analogue of the ψ in that proof. Indeed it can be proven[6] that if there is a small M such that one between p and q is M-definable and the other is finitely satisfiable in M, then $p(x) \otimes q(y) = q(y) \otimes p(x)$. Here definability is a key issue, and *any* attempt to adapt Theorem 1.79 without further hypotheses is bound to fail, as this would prove $p(x) \otimes q(y) = q(y) \otimes p(x)$, and this is generally false: for instance if p is as in Example 2.21, then $x < y \in p(x) \otimes p(y)$ and $y < x \in p(y) \otimes p(x)$.

In fact, in *ordered* structures the only chance that a type has to commute with itself it is that $x = y \in p(x) \otimes p(y)$, and it is easy to see that this implies that p is a realized type: otherwise it would contain all the formulas $x \neq u$ for $u \in \mathfrak{U}$, so $p(x) \otimes p(y)$ contains $x \neq y$, and since there is an order it must then choose between $x < y$ and $y < x$. Clearly $p(y) \otimes p(x)$ will make the opposite choice.

Global types that happen to be simultaneously M-finitely satisfiable and M-definable for a small M, and hence commute with themselves, are of a very special kind. They go under the name of *generically stable*, and the interested reader can consult [Sim15, Chapter 2].

[4] Just ignore formulas in the third free variable.

[5] It is $d_q\varphi$, and not $d_p\varphi$, because as we said the product is backwards.

[6] We will not do it just because it will not be needed later, but the argument is similar to the proof of Theorem 1.79. See [Sim15, Lemma 2.23].

We end this section with a concept that will be exploited in the next one.

Definition 2.32. Let $p \in S_x^{\mathrm{inv}}(\mathfrak{U})$. Set $p^{(1)}(x_0) = p(x_0)$ and define inductively $p^{(n+1)}(x_0, \dots, x_n) = p(x_n) \otimes p^{(n)}(x_0, \dots, x_{n-1})$. We then define $p^{(\omega)} = \bigcup_{n<\omega} p^{(n)}$.

Proposition 2.33. Let $p \in S_x^{\mathrm{inv}}(\mathfrak{U})$. Then any $(b_i)_{i<\omega} \vDash p^{(\omega)} \restriction A$ is A-indiscernible. In particular $\mathrm{EM}((b_i)_{i<\omega}/A) = p^{(\omega)} \restriction A$ is a complete Ehrenfeucht-Mostowski type.

Proof. For all $i_0 < \ldots < i_n$ both $\mathrm{tp}(b_{i_0}, \dots, b_{i_n}/A)$ and $\mathrm{tp}(b_0, \dots, b_n/A)$ are equal to $p^{(n+1)} \restriction A$. □

Definition 2.34. Let $p \in S_x^{\mathrm{inv}}(\mathfrak{U})$. A *Morley sequence of p over A* is any[7] A-indiscernible sequence $(b_i)_{i\in I}$ such that $\mathrm{EM}((b_i)_{i\in I}/A) = p^{(\omega)} \restriction A$.

2.2 Lascar Strong Types

In this section we study a generalization of types that, in turn, will yield a generalization of invariant types. Moreover, this will enable us to prove the basic properties of G_M^{00}.

Bounded Invariant Equivalence Relations

We begin with the following observation: if M is a small model and $a \equiv_M b$, it is not necessarily the case that there is an M-indiscernible sequence starting with $a, b, \dots$

Example 2.35. Let $M \vDash \mathsf{RCF}$, let $a \vDash \mathrm{tp}(+\infty/M)$ and set $b = a+1$. Clearly $b \vDash \mathrm{tp}(+\infty/M)$, but if $a, b, c, \dots$ is the beginning of an M-indiscernible sequence, applying indiscernibility to (a, b) and (b, c) with respect to the formula $x_1 = x_0 + 1$ we find that $c = b + 1 = a + 2$. Applying it to (a, b) and (a, c) yields $c = a + 1$.

So we cannot ask this much. Anyway, in Proposition 2.37 we will show that we can ask a slightly weaker thing. First we will need a lemma.

Lemma 2.36. The binary relation "b, c is the start of an A-indiscernible sequence" is symmetric.

Proof. By the Standard Lemma any A-indiscernible $(a_i)_{i<\omega}$ such that $a_0 = b$ and $a_1 = c$ can be extended to an A-indiscernible $(a_i)_{i\in\mathbb{Z}}$. It is then sufficient to reverse it[8] and cut it before c to have an indiscernible sequence starting with c, b. □

[7] Notice that it can be indexed on any total order I, not necessarily ω.

[8] The reversed sequence will still be A-indiscernible, but may have a different Ehrenfeucht-Mostowski type.

Proposition 2.37. Let M be a small model and suppose that $a \equiv_M b$. Then there is c such that both a, c and c, b start an M-indiscernible sequence.

Proof. Since M is a model, by Proposition 1.70 there is $p \in S^{\mathrm{fs}}(\mathfrak{U}, M)$ which is a coheir of $\mathrm{tp}(a/M) = \mathrm{tp}(b/M)$. Take as $(a_i)_{i<\omega}$ any Morley sequence of p over Mab, which exists because p is M-invariant. To check that the concatenation $a \frown (a_i)_{i<\omega}$ is M-indiscernible it is sufficient to show that for all $i_1 < \ldots < i_n$ we have $\mathrm{tp}(a, a_{i_1}, \ldots a_{i_n}/M) = p^{(n)} \restriction M$, and since $(a_i)_{i<\omega}$ is Ma-indiscernible it suffices to check this for $\mathrm{tp}(a, a_1, \ldots a_n/M)$, i.e. to show that for all $\varphi(x_0, \ldots, x_n) \in L(M)$ we have $\vDash \varphi(a, a_1, \ldots a_n) \leftrightarrow \varphi(a_0, a_1, \ldots a_n)$. First of all notice that, by hypothesis, $\mathrm{tp}(a_0/M) = p \restriction M = \mathrm{tp}(a/M)$, so for $n = 0$ the thesis is true. For $n > 0$, if φ is a counterexample,

$$\varphi(a, x_1, \ldots, x_n) \triangle \varphi(x_0, x_1, \ldots, x_n) \in \mathrm{tp}(a_0, a_1, \ldots, a_n/Ma) = p^{(n)} \restriction Ma$$

So in particular this is a formula in $p^{(n)}$. Let $\tilde{a}_0 \vDash p(x_0)$. By definition and associativity of $\otimes$ we then have

$$\varphi(a, x_1, \ldots, x_n) \triangle \varphi(\tilde{a}_0, x_1, \ldots, x_n) \in p^{(n-1)} \mid \mathfrak{U}\tilde{a}_0$$

Since p is finitely satisfiable in M, this is also true of $p^{(n-1)}$ by Lemma 2.29, and then of $p^{(n-1)} \mid \mathfrak{U}\tilde{a}_0$ by Corollary 2.20. Thus we find $m_1, \ldots, m_n \in M$ such that

$$\vDash \varphi(a, m_1, \ldots, m_n) \triangle \varphi(\tilde{a}_0, m_1, \ldots, m_n)$$

But since $\tilde{a}_0 \equiv_{Ma} a_0$, the formula above is still valid with a_0 in place of $\tilde{a}_0$, and this contradicts $a \equiv_M a_0$.

Now apply the same argument to $b \frown (a_i)_{i<\omega}$, set $c = a_0$ and apply the previous lemma to b, c. □

The fact we just proved will be useful in order to investigate the following concepts, which will in turn be involved in the study of G^{00}:

Definition 2.38. Let α be an ordinal. An equivalence relation E on α-tuples of $\mathfrak{U}$ is called

A-invariant iff for all $a_0 b_0 \equiv_A a_1 b_1$ we have $E(a_0, b_0) \iff E(a_1, b_1)$;

Type-definable over A iff it can be defined with a partial type over A;

Bounded iff it has a small number of equivalence classes, i.e. iff $\mathfrak{U}$ is $|\mathfrak{U}^\alpha/E|^+$-saturated and $|\mathfrak{U}^\alpha/E|^+$-strongly homogeneous.

Proposition 2.39 ([Sim15, Proposition 5.1])**.** Let $A \subseteq \mathfrak{U}$ be small. For an A-invariant equivalence relation on $\mathfrak{U}^\alpha$ the following are equivalent:

1. E is bounded.

2. For all A-indiscernible $(a_i)_{i<\omega}$ with $|a_0| = \alpha$ and all $i, j < \omega$, $E(a_i, a_j)$ holds.
3. For all models $M \supseteq A$ and all $a \equiv_M b$, $E(a, b)$ holds.
4. $|\mathfrak{U}^\alpha / E| \leq 2^{|A|+|T|+|\alpha|}$

Proof.

(1 ⇒ 2) Let $(a_i)_{i<\omega}$ be a counterexample. Since it is A-indiscernible and E is A-invariant, then for *all* $i < j$ we have $\neg E(a_i, a_j)$. By the Standard Lemma we can extend $(a_i)_{i<\omega}$ to a $(a_i)_{i<\kappa}$ for an arbitrarily large κ, thus violating boundedness of E.

(2 ⇒ 3) By Proposition 2.37 and transitivity of E.

(3 ⇒ 4) By Löwenheim-Skolem, take $|M| = |A| + |T|$.

(4 ⇒ 1) We even have an explicit bound. □

Definition 2.40. Suppose that a and b are two α-tuples and that for all A-invariant, bounded equivalence relations E on $\mathfrak{U}^\alpha$ we have $E(a, b)$. Then we say that a and b are *Lascar-equivalent over* A, and write $\mathrm{Lstp}(a/A) = \mathrm{Lstp}(b/A)$ or $a \equiv_{\mathrm{Lstp}_A} b$. Equivalence classes of $\mathrm{Lstp}(-/A)$ are called *Lascar strong types over* A.

Theorem 2.41 ([Sim15, Lemma 5.3])**.** Fix an ordinal α. Then Lascar-equivalence over A on α-tuples has the following characterizations:

1. It is the finest bounded A-invariant equivalence relation on α-tuples.
2. It is the transitive closure of $\Theta_A(a, b)$, the relation "a, b is the start of an A-indiscernible sequence".
3. It is the transitive closure of $\Pi_A(a, b)$, the relation "there is $M \supseteq A$ such that $a \equiv_M b$".

Proof.

1. By definition, $\equiv_{\mathrm{Lstp}_A}$ is the intersection of all A-invariant bounded equivalence relations. It follows immediately that it is A-invariant and, by characterization 2 of Proposition 2.39, it is bounded.
2. We already showed in Lemma 2.36 that Θ_A is symmetric. Its transitive closure Θ^*_A is A-invariant because if a_0c_0, $c_{i-1}c_i$, c_nb_0 all start some A-indiscernible sequence $(d^i_k)_{k<\omega}$ and $f \in \mathrm{Aut}(\mathfrak{U}/A)$ is such that $f(a_0b_0) = a_1b_1$, then $a_1f(c_0)$, $f(c_{i-1})f(c_i)$, $f(c_n)b_1$ start the A-indiscernible sequences $(f(d^i_k))_{k<\omega}$. Moreover Θ^*_A is bounded because it satisfies characterization 2 of Proposition 2.39, hence $a \equiv_{\mathrm{Lstp}_A} b \Rightarrow \Theta^*_A(a, b)$. Again by characterization 2 of Proposition 2.39, $\Theta_A(a, b) \Rightarrow a \equiv_{\mathrm{Lstp}_A} b$, and now it suffices to take transitive closures.

3. Let us show that the transitive closure Π_A^* of Π_A is still A-invariant: indeed $\Pi_A^*(a_0, b_0)$ holds if and only if there is $n \in \omega$, some $(c_i \mid i < n)$ and some $(M_i \supseteq A \mid i < n+1)$, such that $a_0 \equiv_{M_0} c_0$, $c_{i-1} \equiv_{M_i} c_i$ and $c_{n-1} \equiv_{M_n} b_0$. If $a_0 b_0 \equiv_A a_1 b_1$, let $f \in \operatorname{Aut}(\mathfrak{U}/A)$ be such that $f(a_0 b_0) = a_1 b_1$. Then $a_1 \equiv_{f(M_0)} f(c_0)$, $f(c_{i-1}) \equiv_{f(M_i)} f(c_i)$ and $f(c_{n-1}) \equiv_{f(M_n)} b_1$. Moreover Π_A^* is bounded because it satisfies characterization 3 of Proposition 2.39, hence $a \equiv_{\mathrm{Lstp}_A} b \Rightarrow \Pi_A^*(a, b)$. Again by characterization 3 of Proposition 2.39, $\Pi_A(a, b) \Rightarrow a \equiv_{\mathrm{Lstp}_A} b$, and now it suffices to take transitive closures. □

Remark 2.42. Despite the fact that their transitive closures coincide, Example 2.35 shows that, even if M is a model, $\Theta_M \neq \Pi_M$.

These characterizations have a number of consequences, the first one being the following generalization of Proposition 2.37:

Corollary 2.43. Let A be a small set and suppose that $a \equiv_{\mathrm{Lstp}\, A} b$. Then there is $(c_i \mid i < n+1)$ such that a, c_0, $c_{i-1} c_i$ and c_n, b all start an A-indiscernible sequence.

Remark 2.44. Clearly having the same Lstp_A-type implies having the same A-type. If $A = M$ is a model, the converse also holds by Theorem 2.41.

The relation $\equiv_A$ has an "algebraic" counterpart in a subgroup of $\operatorname{Aut}(\mathfrak{U})$, namely $\operatorname{Aut}(\mathfrak{U}/A)$. Another corollary of Theorem 2.41 is that there is such a subgroup for $\equiv_{\mathrm{Lstp}_A}$ too:

Definition 2.45. The subgroup of $\operatorname{Aut}(\mathfrak{U})$ generated by $\bigcup_{M \supseteq A} \operatorname{Aut}(\mathfrak{U}/M)$ is called $\operatorname{Autf}(\mathfrak{U}/A)$.

Corollary 2.46 ([Sim15, Lemma 5.7])**.** $a \equiv_{\mathrm{Lstp}_A} b$ if and only if there is $f \in \operatorname{Autf}(\mathfrak{U}/A)$ such that $f(a) = b$.

Proof. ($\Rightarrow$) By Theorem 2.41 there is some $(c_i \mid i < n)$ and some $(M_i \supseteq A \mid i < n+1)$, such that $a \equiv_{M_0} c_0$, $c_{i-1} \equiv_{M_i} c_i$ and $c_{n-1} \equiv_{M_n} b$. Let $f_i \in \operatorname{Aut}(\mathfrak{U}/M_i)$ be such that $f_0(a) = c_0$, $f_i(c_{i-1}) = c_i$ and $f_n(c_{n-1}) = b$ and set $f = f_n \circ \ldots \circ f_0$.

($\Leftarrow$) Write $f = f_n \circ \ldots \circ f_0$, with $f_i \in \operatorname{Aut}(\mathfrak{U}/M_i)$, and let $f^{(i)} = f_i \circ \ldots \circ f_0$. Then $a \equiv_{M_0} f^{(0)}(a)$, $f^{(i-1)}(a) \equiv_{M_i} f^i(a)$, and $f^{(n-1)}(a) \equiv_{M_n} f^{(n)}(a) = f(a) = b$. Now apply Theorem 2.41. □

As $\equiv_A$ gives rise to the notion of A-invariant types, $\equiv_{\mathrm{Lstp}_A}$, being finer than $\equiv_A$, gives rise to a weaker notion, that we will call Lstp_A-invariance.

Definition 2.47. We call a global type Lstp_A-*invariant* iff it is fixed by $\operatorname{Autf}(\mathfrak{U}/A)$.

The reader may have expected a definition similar in spirit to Definition 2.2. As the following characterization shows, she was right:

Proposition 2.48 ([Sim15, p. 70])**.** Let $p \in S(\mathfrak{U})$. The following are equivalent:

1. For all $f \in \operatorname{Autf}(\mathfrak{U}/A)$ we have $f(p) = p$.
2. For all $\varphi(x; y) \in L$ and $a \equiv_{\mathrm{Lstp}_A} b$ we have $\varphi(x; a) \leftrightarrow \varphi(x; b) \in p$.
3. Every time $d \vDash p$ and $(a_i)_{i<\omega}$ is A-indiscernible, then it is also Ad-indiscernible.

Proof.
(1 ⇒ 2) If $a \equiv_{\mathrm{Lstp}_A} b$ by Corollary 2.46 there is $f \in \operatorname{Autf}(\mathfrak{U}/A)$ such that $f(a) = b$. Then $\varphi(x; a) \in p \iff \varphi(x; b) = \varphi(x; f(a)) \in f(p) = p$.

(2 ⇒ 1) Fix $f \in \operatorname{Autf}(\mathfrak{U}/A)$ and let $\varphi(x; y) \in L$ and $b \in \mathfrak{U}$ be such that $\varphi(x; b) \in p$. If $b = f(a)$, by Corollary 2.46 $a \equiv_{\mathrm{Lstp}_A} b$, so $\varphi(x; a) \in p$. It follows that $\varphi(x; b) = \varphi(x; f(a)) \in f(p)$.

(2 ⇒ 3) Let $\bar{\imath} = i_0 < \ldots < i_n$ and $\bar{\jmath} = j_0 < \ldots < j_n$ and suppose, up to comparing with some $i_n, j_n < k_0 < \ldots < k_n$, that $i_n < j_0$. Then $a_{\bar{\imath}} a_{\bar{\jmath}}$ obviously start an A-indiscernible sequence and they are hence Lstp_A-equivalent. This implies that $d \vDash \varphi(x; a_{\bar{\imath}}) \leftrightarrow \varphi(x; a_{\bar{\jmath}}) \in p$, so $(a_i)_{i<\omega}$ is Ad-indiscernible.

(3 ⇒ 2) Suppose that $\varphi(x; y) \in L(A)$ and $\vDash \varphi(d; a) \mathbin{\triangle} \varphi(d; b)$. Let $(c_i \mid i < n + 1)$ be given by Corollary 2.43 and consider the A-indiscernible sequences starting with $c_i c_{i+1}$, where $c_{-1} = a$ and $c_{n+1} = b$. By hypothesis there must be some $-1 \leq i \leq n$ such that $\vDash \varphi(d; c_i) \mathbin{\triangle} \varphi(d; c_{i+1})$, because otherwise we would have $\vDash \varphi(d; a) \leftrightarrow \varphi(d; b)$. This contradicts the i-th sequence being Ad-indiscernible. □

Remark 2.49. In Remark 2.44 we noticed that having the same Lascar strong type over A implies having the same type over A. Contravariantly, an A-invariant type is Lstp_A-invariant.

Remark 2.50. Even if an Lstp_A-invariant type need not be A-invariant, it will be M-invariant for all models $M \supseteq A$, because by definition $\operatorname{Aut}(M) \subseteq \operatorname{Autf}(A)$.

Remark 2.51. The point raised in Remark 2.7 is clearly still valid if we want to speak of *partial* Lstp_A-invariant types, replacing Aut with Autf. This time too, the simplest thing is to write $p \vdash$ instead of $\in p$.

The Type-Connected Component

Notation 2.52. Until the end of Section 2.2, G will be an $\emptyset$-definable group. We sometimes identify G with $G(\mathfrak{U})$, and G^{00}_A with $G^{00}_A(\mathfrak{U})$.

Lemma 2.53. If H is an A-type-definable subgroup of G of bounded index, then $y^{-1}x \in H$ is an A-invariant, bounded equivalence relation.

Proof. If $a_0b_0 \equiv_A a_1b_1$ then $\mathrm{tp}(b_0^{-1}a_0/A) = \mathrm{tp}(b_1^{-1}a_1/A)$, and since H is A-type-definable this implies $b_0^{-1}a_0 \in H \iff b_1^{-1}a_1 \in H$. Boundedness follows from the fact that the equivalence classes under examination are exactly the cosets of H, which are bounded in number by hypothesis. □

Proposition 2.54. If $M \supseteq A$ is a small model, the projection to the quotient $G(\mathfrak{U}) \to G(\mathfrak{U})/G_A^{00}(\mathfrak{U})$ factors through $S_G(M)$.

Proof. For every A-type-definable $H \supseteq G_A^{00}$ of bounded index, if $a \equiv_M b$ then $b^{-1}a \in H$ by the equivalence of boundedness with point 3 of Proposition 2.39. Since G_A^{00} is defined as the intersection of all such H, the G_A^{00}-coset of $g \in G(\mathfrak{U})$ is decided by $\mathrm{tp}(g/M)$. □

Proposition 2.55. G_A^{00} has bounded index.

Proof. Since Proposition 2.54 is true regardless of the choice of $M \supseteq A$, another application of Proposition 2.39 yields that $y^{-1}x \in G_A^{00}$ is a bounded equivalence relation. □

Proposition 2.56. If M is a small model, G_M^{00} is normal.

Proof. Since conjugating a subgroup does not change its index, it suffices to show that the family of bounded index M-type-definable subgroups of G is stable under conjugacy. This is true because, even if $g \in G(\mathfrak{U}) \setminus G(M)$, by Proposition 2.54 $gG_M^{00}g^{-1}$ only depends on $\mathrm{tp}(g/M)$, and this allows to define it with a type over M. □

The Logic Topology

Definition 2.57. Let $M \mathrel{^+{\supset}} A$ and let $\pi\colon S_G(M) \to G/G_A^{00}$ be the map given by Proposition 2.54. The *M-logic topology* on G/G_A^{00} is defined by declaring C closed iff $\pi^{-1}(C)$ is closed.

In other words, abusing the notation in denoting with π both the projection $G \to G/G_A^{00}$ and the map $S_G(M) \to G/G_A^{00}$ from Proposition 2.54, $C \subseteq G_A^{00}(\mathfrak{U})$ is closed iff $\pi^{-1}(C) \subseteq G(\mathfrak{U})$ is type-definable over M.

Remark 2.58. $\pi\colon S_G(M) \to G/G_A^{00}$ is continuous by definition, and it is surjective inasmuch it lifts a surjective map.

Theorem 2.59 ([Sim15, Lemmas 8.9 and 8.10])**.** G/G_A^{00} is a compact Hausdorff topological group when endowed with the M-logic topology.

Proof. By the previous remark, G/G_A^{00} is a continuous image of a compact space, hence it is compact. To prove Hausdorffness, let $g, h \in G(\mathfrak{U})$ have different G_A^{00}-cosets. Up realizing in $M \mathrel{^+{\supset}} A$ their types over some $M_0 \supseteq A$ such that $|M_0| = |A|$, we may assume without loss of generality that $g, h \in M$. For all $x, y \in G(\mathfrak{U})$ we have $\vDash \big((g^{-1}x \in G_A^{00}) \wedge (h^{-1}y \in G_A^{00})\big) \to$

($y^{-1}x \notin G^{00}_A$) and, since G^{00}_A is type-definable over A, compactness yields $\varphi(x;y) \in L(A)$ such that $\vDash y^{-1}x \in G^{00}_A \to \varphi(x;y)$ and $\vDash \varphi(x;g) \wedge \varphi(y;h) \to y^{-1}x \notin G^{00}_A$. Then $U_g = \{xG^{00}_A \mid\vDash \varphi(x;g)\}$ and $U_h = \{yG^{00}_A \mid\vDash \varphi(y;h)\}$ witness separation of gG^{00}_A and hG^{00}_A and are $L(M)$-open. Hence G^{00}_A is a compact Hausdorff topological *space*. The fact that it is a topological *group* follows from the fact that the group operation and inverse are $\emptyset$-definable: if $\psi(x,y,z) \in L$ defines the function $f(x,y) = x \cdot y^{-1} = z$, then for all $\varphi(z) \in L(M)$ we have $f^{-1}([\varphi(z)]) = [\exists z\ \psi(x,y,z) \wedge \varphi(z)]$. □

Corollary 2.60. There is only one logic topology on G/G^{00}_A.

Proof. Two compact Hausdorff topologies are either incomparable or identical, because as soon as the identity is continuous it is automatically closed. If $M \prec N$ they are comparable, and by amalgamation this suffices. □

This authorizes us to speak of *the* logic topology.

2.3 Forking

In the previous section the notion of an A-invariant type was weakened to that of an Lstp_A-invariant one. We now pursue a further weakening, namely A-non-forking.

Dividing

Definition 2.61. Let $k \in \omega$. A set of formulas is *k-inconsistent* iff all of its subsets of size k are inconsistent.

Definition 2.62 ([TZ12, Definition 7.1.2])**.** Let $\varphi(x;b)$ be an $L(\mathfrak{U})$-formula and $A \subseteq \mathfrak{U}$. We say that φ *divides over A with respect to k* iff there are infinitely many realizations $(b_i)_{i<\omega}$ of $\mathrm{tp}(b/A)$ such that $\{\varphi(x;b_i) \mid i < \omega\}$ is k-inconsistent. If we simply say that $\varphi(x;b)$ *divides over* A we mean that this happens for some k. *A-dividing* is the same thing as dividing over A.

Remark 2.63. In the previous definition it makes no difference if $\varphi(x;y) \in L$ or $\varphi(x;y) \in L(A)$: if b contains parameters from A they will not change in any $b_i \equiv_A b$, therefore we may safely hide them.

In other words a formula divides over A if there is a $k \in \omega$ such that the set it defines is "small enough" that we can find infinitely many "A-indiscernible copies" of it and arrange them in a way that no k of them intersect.

Example 2.64 ([TZ12, p. 108])**.** Here is a formula in DLO that divides over $\emptyset$. Let $b = (c_0, c_1)$ and consider $\varphi(x;b) = c_0 < x < c_1$. Extend c_0, c_1 to a strictly increasing sequence $(c_i)_{i<\omega}$ and set $b_i = (c_{2i}, c_{2i+1})$. Then $\{\varphi(x;b_i) \mid i < \omega\} = \{c_{2i} < x < c_{2i+1} \mid i < \omega\}$ is 2-inconsistent. A non-example of dividing formula is $x > b$.

Remark 2.65. If $A \subseteq B$ and φ divides over B, it also divides over A, since the type of a tuple over B decides its type over A.

We would like to be able to treat dividing as a "notion of smallness", in some suitable sense; for example we could hope for A-dividing formulas to form an ideal, i.e. for their negation to form a filter. In this direction, the following lemma seems (deceptively, as we will see) promising.

Lemma 2.66 ([Sim15, Lemma 5.13])**.** If $\vDash \varphi(x;b) \to \psi(x;c)$ and ψ divides over A with respect to k, then φ does too.

Proof. Let $(c_i)_{i<\omega}$ witness that φ divides over A. Fix, for all $i \in \omega$, some $f_i \in \mathrm{Aut}(\mathfrak{U}/A)$ such that $f_i(c) = c_i$, and then set $b_i = f_i(b)$. Then clearly $b_i \equiv_A b$ and $\vDash \varphi(x;b_i) \to \psi(x;c_i)$. Then any $\{\varphi(x;b_{i_j}) \mid j < k\}$ implies $\{\varphi(x;c_{i_j}) \mid j < k\}$, which is inconsistent by hypothesis. □

In addition to providing "downward closure" of dividing, the previous lemma also implies that this notion is stable under equivalence.

Corollary 2.67. If $\varphi(x;b)$ defines the same set as $\psi(x;c)$ then φ divides over A if and only if ψ does.

The following proposition provides an equivalent definition of dividing.

Proposition 2.68 ([TZ12, Lemma 7.1.4])**.** $\varphi(x;b)$ divides over A if and only if it we can choose the witnessing sequence $(b_i)_{i<\omega}$ to be A-indiscernible and to start with $b_0 = b$.

Proof. If $(b_i)_{i<\omega}$ is not A-indiscernible, apply the Standard Lemma to find $(\tilde{b}_i)_{i<\omega}$ that is and such that $\mathrm{EM}(\tilde{b}_i)_{i<\omega} \supseteq \mathrm{EM}(b_i)_{i<\omega}$. Notice that the latter object knows about k-inconsistency of $\{\varphi(x;y_i) \mid i < \omega\}$. To complete the proof, just find $f \in \mathrm{Aut}(\mathfrak{U}/A)$ such that $f(\tilde{b}_0) = b$ and consider $(f(\tilde{b}_i))_{i<\omega}$. The other implication is trivial. □

Permanent Assumption 2.69. Proposition 2.68 will be used without mention throughout the rest of the present thesis.

Definition 2.70. A partial type $\pi(x)$ *divides over* A iff it implies a formula that divides over A.

Proposition 2.71 ([Sim15, Proposition 5.14 (1)])**.** Consistent partial A-types do not A-divide.

Proof. Saying that $\pi(x) \vdash \varphi(x;b)$ means that there is $\psi(x) \in L(A)$, a conjunction of formulas from $\pi(x)$, such that $\vDash \forall x\ \psi(x) \to \varphi(x;b)$. Since $\psi \in L(A)$, the same remains true replacing b with any $b_i \equiv_A b$. This means that for all A-indiscernible $(b_i)_{i<\omega}$ starting with $b = b_0$, for all i we have $\pi(x) \vdash \varphi(x;b_i)$. Therefore $\{\varphi(x;b_i) \mid i < \omega\}$ being inconsistent would contradict consistency of $\pi(x)$. □

Lemma 2.72 ([Sim15, Lemma 5.17])**.** Let $\pi(x)$ be a partial type over Ab. Then π does not A-divide if and only if for every A-indiscernible $(b_i)_{i<\omega}$ starting with b there is $a \vDash \pi$ such that $(b_i)_{i<\omega}$ is Aa-indiscernible.

Proof. ($\Rightarrow$) Hiding parameters coming from A, write $\pi(x) = \pi(x; b)$. Since π does not divide, $\bigcup_{i<\omega} \pi(x; b_i)$ must be consistent, hence have a realization $\tilde{a}$. By the Standard Lemma, let $(\tilde{b}_i)_{i<\omega}$ be an $A\tilde{a}$-indiscernible sequence such that $\mathrm{EM}((\tilde{b}_i)_{i<\omega}/A\tilde{a}) \supseteq \mathrm{EM}((b_i)_{i<\omega}/A\tilde{a})$. Since the restriction of the latter to A is complete (because $(b_i)_{i<\omega}$ is A-indiscernible) there is $f \in \operatorname{Aut}(\mathfrak{U}/A)$ such that $f(\tilde{b}_i) = b_i$. It is then sufficient to take $a = f(\tilde{a})$.

($\Leftarrow$) By Lemma 2.66 and compactness it suffices to show that no finite conjunction $\varphi(x; b)$ of formulas in π divides. If $(b_i)_{i<\omega}$ is a candidate to witness dividing, let a be given by the hypotheses. Then $\vDash \varphi(a; b)$, and by Aa-indiscernibility also $\vDash \varphi(a; b_i)$. Hence $\{\varphi(x; b_i) \mid i < \omega\}$ is consistent. □

Forking

Unfortunately, it is *not* generally true that A-dividing formulas form an ideal. The remedy for this is the obvious one, i.e. to consider the ideal generated, which by Lemma 2.66 amounts to nothing else that allowing disjunctions. We state the definition directly for types.

Definition 2.73. A partial type *forks over A* iff it implies a disjunction $\bigvee_{j<n} \varphi_j(x; b^j)$ such that each $\varphi_j(x; b^j)$ divides over A. Saying that a partial type "is A-non-forking" means that it does not fork over A. If we say that a formula φ forks over A, we mean that $\{\varphi\}$ forks over A.

Remark 2.74. A-non-forking types form a closed subset of $S_x(\mathfrak{U})$, since by compactness if a type forks over A then a finite subset of it already does, and this yields an open subset consisting of types forking over A.

Remark 2.75. If $A \subseteq B$ and a type divides over B, it also divides over A.

For an example of a type that forks but does not divide see [Sim15, Example 5.15]. Here is a first case where forking collapses onto diving:

Proposition 2.76 ([Sim15, Proposition 5.14 (3)])**.** If $A \subset^+ M$ and $p \in S(M)$, then p forks over A if and only if p divides over A. As a special case, for global types[9] forking over a small A is the same as A-dividing.

Proof. Let[10] $p(x) \vdash \bigvee_{i<n} \varphi_i(x; b)$. Since the $\varphi_i(x; b)$ need not be $L(M)$-formulas, we cannot take for granted that one of them belongs to p; in other words, even if p is complete *as a type over M*, it is *partial* when regarded as a global type. This is where saturation enters the picture.

[9]If we speak of global types, the $(b_i)_{i<\omega}$ may live outside $\mathfrak{U}$.

[10]Since, for instance, it does not matter if we call $\varphi(x; a, b, c, d, e, f, g)$ the formula $x = a$, up to adding unused parameters to $\varphi_i(x, b^i)$ we can fix the same b for all φ_i.

By compactness there is some finite $C \subseteq M$ such that $p(x) \upharpoonright C \vdash \bigvee_{i<n} \varphi_i(x; b)$. By saturation, find *inside* M some $\tilde{b} \equiv_{C \cup A} b$. Then $p(x) \upharpoonright C \vdash \bigvee_{i<n} \varphi_i(x; \tilde{b})$, because $\tilde{b} \equiv_C b$. Since we are now dealing with a disjunction of $L(M)$-formulas, there is i such that $\varphi_i(x; \tilde{b}) \in p(x)$, and this formula divides over A because $\tilde{b} \equiv_A b$. □

Considering non-forking types instead of non-dividing ones pays off in the following way:

Proposition 2.77 ([Sim15, Proposition 5.14 (4)])**.** If $\pi(x)$ is a partial type over $B \supseteq A$, then π does not fork over A if and only if it can be extended to a $p \in S_x(B)$ that does not fork over A.

Proof. Since, by compactness and Lemma 2.66, it suffices to check formulas in $L(B)$, such a p exists if an only if this set of formulas is consistent:

$$\pi(x) \cup \{\neg\varphi(x; b) \mid \varphi(x; b) \in L(B) \text{ forks over } A\}$$

If it is not, then by compactness there are finitely many A-forking φ_i such that $\pi(x) \wedge \bigwedge_{i<n} \neg\varphi_i(x; b) \vdash \perp$, or in other words $\pi(x) \vdash \bigvee_{i<n} \varphi_i(x; b)$. Since forking is closed under disjunctions, $\pi(x)$ forks over A. The other implication is trivial. □

Proposition 2.78 ([Sim15, Proposition 5.14 (6)])**.** Let A be small. If $p \in S_x(\mathfrak{U})$ is Lstp_A-invariant, then it is A-non-forking.

Proof. Since A is small it is sufficient to prove non-dividing. Suppose $p(x) \vdash \varphi(x; b)$ and $(b_i)_{i<\omega}$ is an A-indiscernible sequence starting with b. Clearly[11] any two b_i, b_j start an indiscernible sequence, hence by Theorem 2.41 they have the same Lascar strong type over A. Then, by Proposition 2.48, $p(x) \vdash \varphi(x; b_i)$ for all $i < \omega$, so $\{\varphi(x; b_i) \mid i < \omega\}$ cannot be inconsistent because p is not. □

Corollary 2.79. Global M-invariant types do not fork over M.

Proof. By the previous proposition and Remark 2.50. □

Back to Example 2.18, types at ∞ are $\emptyset$-invariant, hence for all small M they do not fork over M: indeed, in this case they do not even fork over $\emptyset$: it is sufficient to check dividing and it is easy to see that half-lines cannot divide.

The type $p(x) = \mathrm{tp}(\mathbb{Q}^+/\mathfrak{U})$ is finitely satisfiable in[12] $\mathbb{Z}$, hence $\mathbb{Z}$-invariant; let us see how it fails to divide over $\mathbb{Z}$. Take a formula in $p(x)$, without loss of generality of the form $a < x < b$. We must have $b > \mathbb{Z}$, and there ought

[11] Just discard some elements from $(b_i)_{i<\omega}$ and apply Lemma 2.36 if needed.

[12] Obviously, $\mathbb{Z}$ is not a model of DLO. If the reader prefers, she can substitute $\mathbb{Z}$ with $\mathbb{Q}$ in the whole example.

to be some point of $\mathbb{Z}$, say 17, such that $a < 17$. Then, for any indiscernible $(a_i b_i)_{i<\omega}$, we have that each b_i must still be greater than $\mathbb{Z}$, and a_i must still be less than 17. Then 17 witnesses that $a < x < b$ does not divide.

On the other hand, a non-invariant global type in one variable must divide over any small model M: as we saw in Example 2.18 we can find $a < b$ such that $[a, b] \cap M = \emptyset$, and then the formula $a < x < b$ divides over M. To see this, consider the "hole in M" where a and b lie, i.e.

$$H = \{u \in \mathfrak{U} \mid \forall m \in M\ m < a \to m < u\} \cap \{u \in \mathfrak{U} \mid \forall m \in M\ m > b \to m > u\}$$

By definition, any $a_i b_i \in H$ such that $a_i < b_i$ realizes $\operatorname{tp}(ab/M)$. It suffices then to find $\aleph_0$ pairwise disjoint segments (a_i, b_i) inside H.

Lemma 2.80. Suppose that $A \subseteq B$ and $\operatorname{tp}(a/B)$ does not fork over A. Then there is a small model $M \mathrel{^{+}{\supset}} B$ such that $\operatorname{tp}(a/M)$ does not fork over A.

Proof. Since $\operatorname{tp}(a/B)$ does not fork over A, by Proposition 2.77 it has a global A-non-forking extension; let q be its restriction to a small $M_0 \mathrel{^{+}{\supset}} B$. Clearly q does not fork over A and $q \upharpoonright B = \operatorname{tp}(a/B)$. Let $a_0 \vDash q$ and take $f \in \operatorname{Aut}(\mathfrak{U}/B)$ such that $f(a_0) = a$. If we set $M = f(M_0)$, then since $q = \operatorname{tp}(a_0/M_0)$ does not fork over A, neither $f(q) = \operatorname{tp}(f(a_0)/f(M_0)) = \operatorname{tp}(a/M)$ does. □

Independence

Notation 2.81. $a \mathop{\smile\!\!\!\!\!\!|}_A b$ is a shorthand for "$\operatorname{tp}(a/Ab)$ is A-non-forking".

Lemma 2.82 ([Sim15, Lemma 5.18])**.** $\mathop{\smile\!\!\!\!\!\!|}$ is left-transitive, i.e. if $A \subseteq B$, $a \mathop{\smile\!\!\!\!\!\!|}_A B$, and $b \mathop{\smile\!\!\!\!\!\!|}_{Aa} Ba$, then $ab \mathop{\smile\!\!\!\!\!\!|}_A B$.

Proof. By Lemma 2.80 there is $\tilde{M} \mathrel{^{+}{\supset}} B$ such that $\operatorname{tp}(a/\tilde{M})$ does not fork over B. Since $\operatorname{tp}(b/Ba)$ does not fork over Aa, it has a global Aa-non-forking extension q. If $\tilde{b} \vDash q \upharpoonright \tilde{M}a$ then there is $f \in \operatorname{Aut}(\mathfrak{U}/Ba)$ such that $f(\tilde{b}) = b$, and setting $M = f(\tilde{M})$ we have that $\operatorname{tp}(b/Ma)$ does not fork over Aa and that $\operatorname{tp}(a/M)$ does not fork over A. By Proposition 2.76 it is sufficient to show that $\operatorname{tp}(a, b/M)$ does not A-divide. Suppose that $\varphi(x, y; c) \in L(M)$ and $(c_i)_{i<\omega}$ witnesses A-dividing. Since $c \in M$, and $\operatorname{tp}(a/M)$ does not fork over A, by Lemma 2.72 there is $f \in \operatorname{Aut}(\mathfrak{U}/Ac)$ such that $(c_i)_{i<\omega}$ is $Af(a)$-indiscernible. Since $\operatorname{tp}(b/Ma)$ does not fork over Aa, then $\operatorname{tp}(f(b)/Acf(a))$ does not fork over $Af(a)$, and as a special case it does not $Af(a)$-divide. Again by Lemma 2.72 we can find $g \in \operatorname{Aut}(\mathfrak{U}/Af(a)c)$ such that $(c_i)_{i<\omega}$ is $Ag(f(a))g(f(b))$-indiscernible[13]. Since $\vDash \varphi(a, b; c)$ and $f(c) = c = g(c)$, this implies that $\vDash \varphi(g(f(a)), g(f(b)); c_i)$, so $\{\varphi(x, y; c_i) \mid i < \omega\}$ is consistent. □

[13]Of course $g(f(a)) = f(a)$, however we feel that writing $g(f(a))$ improves readability.

Corollary 2.83 ([Sim15, Corollary 5.20])**.** If $p(x), q(y) \in S^{\mathrm{inv}}(\mathfrak{U})$ are A-non-forking, then $p(x) \otimes q(x)$ is too.

Proof. If $p(x) \otimes q(x)$ forks over A, then there is a small $B \supseteq A$ such that $p(x) \otimes q(x) \upharpoonright B$ does. Up to enlarging B, we may assume $p, q \in S^{\mathrm{inv}}(\mathfrak{U}, B)$. Then, if $b \vDash q \upharpoonright B$ and $a \vDash p \upharpoonright Bb$, by Corollary 2.26 $\mathrm{tp}(a, b/B) = p \otimes q \upharpoonright B$. Since nor p nor q fork over A, neither their restrictions do, therefore $b \mathop{\smile\!\!\!\!\!\!\mid}_A B$ and $a \mathop{\smile\!\!\!\!\!\!\mid}_A Bb$, and a fortiori $a \mathop{\smile\!\!\!\!\!\!\mid}_{Ab} Bb$. Thus, by left-transitivity, $\mathrm{tp}(a, b/B)$ does not fork over A. □

Lemma 2.84. If $(c_j)_{j<\omega}$ is a Morley sequence of $p \in S_x^{\mathrm{inv}}(\mathfrak{U}, A)$ over A, then for all $j \in \omega$, we have[14] $c_{>j} \mathop{\smile\!\!\!\!\!\!\mid}_A c_j$.

Proof. Since cutting off the first $j-1$ elements of $(c_j)_{j<\omega}$ still results in a Morley sequence of p over A we can assume $j = 0$. By compactness it is enough to set $\bar{c} = (c_1, \dots, c_k)$ and check that $\mathrm{tp}(\bar{c}/Ac_0)$ does not fork over A. Since p is A-invariant, by Corollary 2.26 we have $\mathrm{tp}(\bar{c}/Ac_0) = p^{(k-1)} \upharpoonright Ac_0$. Since A-invariance implies A-non-forking and these notions are preserved under products, $\mathrm{tp}(\bar{c}/Ac_0)$ does not fork over A. □

We now prove some technical results that will be needed later.

Lemma 2.85. Let $b \mathop{\smile\!\!\!\!\!\!\mid}_A c$ and let $(c_j)_{j<\omega}$ be a Morley sequence of p over A with $c_0 = c$. Then there is $(d_j)_{j<\omega}$, a Morley sequence of p over A with $d_0 = c$, such that $(c_j)_{j<\omega} \equiv_{\mathrm{Lstp}_A} (d_j)_{j<\omega}$ and $b \mathop{\smile\!\!\!\!\!\!\mid}_A (d_j)_{j<\omega}$.

Proof. By Lemma 2.80 there is $M \supseteq Ac$ such that $\mathrm{tp}(b/M)$ does not fork over A. Let $N \supseteq M(c_j)_{j<\omega}$ be a small model and let $q \in S(N)$ be an A-non-forking extension of $\mathrm{tp}(b/M)$. Let $\tilde{b} \vDash q$. Since $\mathrm{tp}(\tilde{b}/A(c_j)_{j<\omega})$ is implied by q it cannot fork over A. Now let $f \in \mathrm{Aut}(\mathfrak{U}/M)$ be such that $f(\tilde{b}) = b$ and set $d_j = f(c_j)$. Since $c \in M$, we have $d_0 = c$. Moreover $(c_j)_{j<\omega} \equiv_M (d_j)_{j<\omega}$ and $A \subseteq M$; this implies that $(c_j)_{j<\omega} \equiv_{\mathrm{Lstp}_A} (d_j)_{j<\omega}$ and that $(d_j)_{j<\omega}$ is still a Morley sequence of p over A. To complete the proof, notice that $\mathrm{tp}(b/A(d_j)_{j<\omega}) = \mathrm{tp}(f(\tilde{b})/A(f(c_j))_{j<\omega}) = \mathrm{tp}(\tilde{b}/A(c_j)_{j<\omega})$, so it does not fork over A. □

Definition 2.86. An *extension base* is a set A such that no type in $S(A)$ forks over A.

Proposition 2.87. Models are extension bases.

Proof. Suppose $S_x(M) \ni p(x) \vdash \bigvee_{i<n} \varphi_i(x; c)$, each φ_i dividing over M and $\varphi_i(x; y) \in L(M)$. By compactness, there is an $L(M)$-formula $\psi(x) \in p(x)$ such that $\vDash \psi(x) \to \bigvee_{i<n} \varphi_i(x; c)$. Since p is a type over M, and M is a model, there is $m \in M$ such that $\vDash \psi(m)$ and — say — $\vDash \varphi_0(m; c)$. Then no M-indiscernible $(c_j)_{j<\omega}$ starting with c can witness dividing since, for all $j \in \omega$, by M-indiscernibility $\vDash \varphi_0(m; c_j) \leftrightarrow \varphi_0(m; c_0)$. □

[14] $c_{>j}$ means $\{c_k \mid k > j\}$.

The fact that forking formulas can be thought as the "small" ones can sometimes mean more than simply forming an ideal. For instance, in ω-stable theories, where Morley rank is defined (see for example [Mar02, Chapter 6]), the forking extensions of a type are exactly the ones with smaller Morley rank.

Forking was born in the study of, and behaves best in, stable theories. The symmetry in the symbol $\mathop{\smile\!\!\!\!\!\mid}$ is probably due to the fact that, in stable theories, $a \mathop{\smile\!\!\!\!\!\mid}_A b \iff b \mathop{\smile\!\!\!\!\!\mid}_A a$. This is *not* true in general, and for instance it fails in DLO as [Sim15, Example 5.24] shows. A class where $\mathop{\smile\!\!\!\!\!\mid}$ behave nicely in this respect is the one of *simple* theories, treated for example in [Cas11, Wag00, TZ12]. DLO is not simple, but the random graph is. The relation $\mathop{\smile\!\!\!\!\!\mid}$ is sometimes called *forking independence*.

2.4 Keisler Measures

Ultrafilters on X can be seen as finitely additive probability measures, defined on the whole $\mathscr{P}(X)$, that are only allowed to take values in $\{0,1\}$. Similarly, types in $S_x(A)$ are $\{0,1\}$-valued finitely additive measures, defined on the Boolean algebra[15] $\operatorname{Def}_x(A)$. From this viewpoint, measures are nothing but "fuzzy" types, i.e. "types with values in the real interval $[0,1]$". As we will see, simply replacing braces with brackets gives rise to a very rich topic, especially in the NIP context. The purpose of this section is to give definitions and collect some basic results that hold without any hypothesis on T.

Measures on Definable Sets

Definition 2.88. Let $A \subseteq \mathfrak{U}$. A *Keisler measure* over A is a finitely additive probability measure $\mu\colon \operatorname{Def}_x(A) \to [0,1]$.

Notation 2.89. Variables will be often specified near to the measure, i.e. if $\mu\colon \operatorname{Def}_x(A) \to [0,1]$ we will often write "$\mu(x)$ is a measure over A". Another notation that will be used is μ_x, and it means the same thing as $\mu(x)$. The space of all Keisler measures over A in variables x is denoted $\mathfrak{M}_x(A)$. If $B \subseteq A$ we will write $\mu_x \restriction B$ for $\mu_x \restriction \operatorname{Def}_x(B)$. Clearly, if $|x| = |y|$, there is a canonical isomorphism between $\operatorname{Def}_x(A)$ and $\operatorname{Def}_y(A)$, and it induces an homeomorphism between $\mathfrak{M}_x(A)$ and $\mathfrak{M}_y(A)$; if we have a measure $\mu(x)$ (or μ_x), we will denote $\mu(y)$ (or μ_y) its image under the previous identification.

Proposition 2.90. $\mathfrak{M}_x(A)$ is a closed subset of $[0,1]^{\operatorname{Def}_x(A)}$, hence it is compact Hausdorff.

Proof. Keisler measures over A are those $[0,1]$-valued functions that satisfy $\mu(x = x) = 1$, $\mu(\neg\varphi(x)) = 1 - \mu(\varphi(x))$ and $\mu(\varphi(x) \wedge \psi(x)) + \mu(\varphi(x) \vee \psi(x)) =$

[15] Since we are only talking of finite additivity, it is not necessary to work on a σ-algebra.

$\mu(\varphi(x)) + \mu(\psi(x))$ for all $\varphi, \psi \in L(A)$. These three conditions are closed in the product topology. □

Example 2.91 ([Sim15, Example 7.2])**.** To every $p \in S_x(A)$ we can associate an element of $\mathfrak{M}_x(A)$, still denoted p, setting $p(\varphi(x)) = 1$ if $\varphi(x) \in p$ and $p(\varphi(x)) = 0$ if $\neg\varphi(x) \in p$.

Before any other example we notice that the type space embeds nicely in the space of measures and explain the ambivalent nature of Keisler measures.

Proposition 2.92. *$S_x(A)$ is closed in $\mathfrak{M}_x(A)$, and the usual topology on $S_x(A)$ coincides with the topology it inherits as a subspace of $\mathfrak{M}_x(A)$. In other words the inclusion map $S_x(A) \hookrightarrow \mathfrak{M}_x(A)$ is an homeomorphism with its image.*

Proof. All point-wise limits of functions with values in $\{0, 1\}$ still have values in $\{0, 1\}$, and it is easy to see that a point-wise limit of types is still a type. Moreover, if $p \in S_x(A)$, then saying $\varphi \in p$, $p(\varphi) = 1$ or $p(\varphi) > 1/2$ makes no difference. This shows that every $[\varphi]$ can be seen as an open subset in the product topology, and if a compact Hausdorff topology is finer than another compact Hausdorff one, then they are equal, since as soon as the identity is continuous it is automatically closed. □

Regular Measures on Types

Now, $\mathrm{Def}_x(A)$ need not be a σ-algebra, since the intersection of infinitely many definable sets need not be definable. Anyway, such a thing is precisely the same as a partial type. We are therefore led to see $\mathrm{Def}_x(A)$ as the Boolean algebra of the clopen sets of $S_x(A)$ and try to extend Keisler measures to real σ-additive measures defined on all the Borel subsets of the type space. Such an extension will be unique if we impose a constraint, namely regularity.

Definition 2.93. Let X be a topological space. A measure μ on the Borel σ-algebra of X is *regular* iff for all Borel $B \subseteq X$

$$\sup_{\substack{C \subseteq B \\ C \text{ compact}}} \mu(C) = \inf_{\substack{U \supseteq B \\ U \text{ open}}} \mu(U)$$

Theorem 2.94 ([Sim15, p. 99])**.** *Each finitely additive probability measure on definable sets extends to a unique regular Borel σ-additive probability measure on types.*

Proof. Let $\mu \in \mathfrak{M}_x(A)$, and identify $\mathrm{Def}_x(A)$ with the Boolean algebra of clopen subsets of $S_x(A)$. Thus $\mu([\varphi]) = \mu(\varphi)$. We first extend μ to open sets approximating from the inside, and to closed sets approximating from the outside, i.e. if U is open and C is closed we set

$$\mu(U) = \sup_{\substack{[\varphi] \subseteq U \\ \varphi(x) \in L(A)}} \mu(\varphi(x)) \qquad \mu(C) = \inf_{\substack{[\varphi] \supseteq C \\ \varphi(x) \in L(A)}} \mu(\varphi(x))$$

Claim. If B is closed or open, then

$$\sup_{\substack{C\subseteq B\\ C \text{ closed}}} \mu(C) = \inf_{\substack{U\supseteq B\\ U \text{ open}}} \mu(U) = \mu(B) \tag{2.1}$$

Proof of the Claim. Up to considering $B^{\complement}$, it is sufficient to show it for B open. If $C \subseteq B$ is closed and $B = \bigcup_{i\in I}[\varphi_i]$, then by compactness there is a finite $F \subseteq I$ such that $C \subseteq \bigcup_{i\in F}[\varphi_i] = \left[\bigwedge_{i\in F}\varphi_i\right] \subseteq B$. This shows that

$$\sup_{\substack{C\subseteq B\\ C \text{ closed}}} \mu(C) = \sup_{\substack{[\varphi]\subseteq B\\ \varphi(x)\in L(A)}} \mu(\varphi(x)) = \mu(B) = \inf_{\substack{U\supseteq B\\ U \text{ open}}} \mu(U)$$

□ CLAIM

Now we want to show that if U is open and $C \subseteq U$ is closed, then $\mu(U \setminus C) = \mu(U) - \mu(C)$. Notice that μ is sub-additive on open sets by definition, and by the Claim for all $\varepsilon > 0$ there is an open $V \supseteq C$ such that $\mu(V) \leq \mu(C) + \varepsilon$. By subadditivity

$$\mu(U) \leq \mu(U \cup V) \leq \mu(U \setminus C) + \mu(V) \leq \mu(U \setminus C) + \mu(C) + \varepsilon$$

Since ε was arbitrary we have $\mu(U) \leq \mu(U \setminus C) + \mu(C)$. As for the other inequality, applying the Claim again yields

$$\mu(C) + \mu(U \setminus C) = \mu(C) + \sup_{\substack{D_0\subseteq U\setminus C\\ D_0 \text{ closed}}} \mu(D_0) \leq \sup_{\substack{D\subseteq U\\ D \text{ closed}}} \mu(D) = \mu(U)$$

Where the inequality is due to the fact that each $C \cup D_0$ counts as a D.

Now, if there has to be any hope of extending μ to a *regular* measure on the the Borel σ-algebra of $S_x(A)$, then the first equality in (2.1) must hold for all Borel B. If we manage to prove this, then we can take (2.1) as the definition of μ on Borel sets, and uniqueness will come for free. All we have to do, then, is to prove that the family of subsets of $S_x(A)$ satisfying the first equality in (2.1) is a σ-algebra. Notice that stability under complement is trivial. Countable unions will require some calculations.

Let $B = \bigcup_{i<\omega} B_n$, where each B_n satisfies (2.1), and fix $\varepsilon > 0$. If we find an open U and a closed C such that $C \subseteq B \subseteq U$ and $\mu(U) - \mu(C) < \varepsilon$ we are done. By hypothesis, for all $n \in \omega$ we can choose a closed C_n and an open U_n such that $C_n \subseteq B_n \subseteq U_n$ and $\mu(U_n \setminus C_n) < \varepsilon/2^{n+2}$. Let $V_n = \bigcup_{k<n} U_k$, $U = \bigcup_{n\in\omega} U_n$ and $C = \bigcup_{n\in\omega} C_n$. By compactness, if $[\varphi] \subseteq U$, then $[\varphi] \subseteq V_n$ for some n, and this implies that $\mu(U) = \lim_{n\to\infty} \mu(V_n)$. Hence there is some $N \in \omega$ such that for all $n > N$ we have $\mu(U) - \mu(V_n) < \varepsilon/2$. Therefore

$$\mu(U) - \mu(C) = \lim_{n\to\infty} \mu(V_n) - \mu(C) < \mu(V_N) - \mu(C) + \frac{\varepsilon}{2} \leq \mu(V_N \setminus C) + \frac{\varepsilon}{2}$$
$$\leq \sum_{n<N} \mu(U_n \setminus C) + \frac{\varepsilon}{2} \leq \sum_{n<N} \mu(U_n \setminus C_n) + \frac{\varepsilon}{2} = \Big(\sum_{n<N} \mu(U_n) - \mu(C_n)\Big) + \frac{\varepsilon}{2}$$
$$< \varepsilon$$

□

Notice that, given a regular Borel probability μ on $S_x(A)$, one can recover a finitely additive probability on $\mathrm{Def}_x(A)$ simply restricting μ to clopen sets, and the unique regular extension of this measure to $S_x(A)$ will obviously be μ. We therefore blur the distinction between the two concepts:

Notation 2.95. A Keisler measure μ_x will be seen simultaneously as a finitely additive probability measure on $\mathrm{Def}_x(A)$ and as a σ-additive, regular probability measure on the Borel subsets of $S_x(A)$. When we say *measure* we mean Keisler measure, unless specified otherwise. A measure μ_x will be called *global* if it is a measure on $S_x(\mathfrak{U})$.

Example 2.96 ([Sim15, Example 7.2]). A little calculation shows that $\mathfrak{M}_x(A)$ is convex, i.e. the average of measures is still a measure. For instance given, for $i < n$, some $p_i \in S_x(A)$ and some $a_i \in [0,1]$ such that $\sum_{i<n} a_i = 1$, we can define the measure $\sum_{i<n} a_i p_i(x)$. Of course this need not be a type, i.e. $S_x(A)$ is not convex. For instance, in DLO, the measure $\frac{1}{3}p_{-\infty}(x) + \frac{2}{3}p_{+\infty}(x)$, gives measure 1/3 (resp. 2/3) to left(resp. right)-unbounded, right(resp. left)-bounded definable subsets, 0 to the bounded ones, and 1 to definable subsets which are unbounded on both sides.

Example 2.97 ([Sim15, Example 7.2]). In RCF, any Borel probability measure P on $\mathbb{R}$ corresponds to a Keisler measure on $S_x(\mathfrak{U})$ given by $\mu(\varphi(x)) = P(\varphi(\mathbb{R}))$.

Example 2.98 ([Sim15, Example 7.2]). Each finite structure[16] can be equipped with the normalized counting measure. If $\mathcal{U} \in \beta\omega$ and $(M_n)_{n<\omega}$ is a sequence of finite L-structures and μ_n is such a measure relative to M_n, then $\mu(\varphi(x; [b_n]_\mathcal{U})) = \lim_{n\to\mathcal{U}} \mu_n(\varphi(x; b_n))$ is a measure over the ultraproduct $\prod_\mathcal{U} M_n$.

Support and Properties of Measures

Definition 2.99. If $\mu \in \mathfrak{M}_x(A)$, its *support* $S(\mu)$ consists of the $p \in S_x(A)$ such that if $\varphi \in p$ then $\mu(\varphi) > 0$. If $p \in S(\mu)$ we also say that *p is weakly random for μ.*

In other words, thinking of μ as a measure on $S_x(A)$, its support is made of the points that have no measure zero neighbourhood.

Remark 2.100. $S(\mu)$ is a closed set, since it can be written as $\bigcap_{\mu(\varphi)=0}[\neg\varphi]$.

Lemma 2.101. If $\mu(\varphi(x)) > 0$, then there is $p \in S(\mu)$ such that $\varphi \in p$. In particular, $\mu(S(\mu)) = 1$.

Proof. All we have to do is to find some $p \in [\varphi] \cap \bigcap_{\mu(\psi)=0}[\neg\psi]$, and this is non-empty by compactness and finite additivity. The second statement follows, since then $\mu((S(\mu))^\complement) = \sup_{[\psi]\subseteq(S(\mu))^\complement} \mu(\psi) = \sup_{\mu(\psi)=0} 0 = 0$. □

[16] Of course, in this case we drop the hypothesis that its theory has infinite models.

Lemma 2.102. Let B be a Borel subset of $S_x(\mathfrak{U})$. Then the space $\mathfrak{M}_x^B(\mathfrak{U})$ of measures such that $S(\mu) \subseteq B$ is homeomorphic to the space $\mathfrak{M}(B)$ of measures on B.

Proof. Define $f \colon \mathfrak{M}_x^B(\mathfrak{U}) \to \mathfrak{M}(B)$ as $f\mu(X \cap B) = \mu(X)/\mu(B)$. A closer look shows that this is simply $\mu(X)$, since $\mu(B) \geq \mu(S(\mu)) = 1$. This does not depend on X, i.e. if $X \cap B = Y \cap B$ then $f\mu(X \cap B) = f\mu(Y \cap B)$, because $\mu(X \cap B^{\complement}) \leq \mu(B^{\complement}) \leq \mu((S(\mu))^{\complement}) = 0$, and similarly $\mu(Y \cap B^{\complement}) = 0$. Moreover, as an obvious consequence of the fact that $f\mu(X \cap B) = \mu(X)$, for all open $U \subseteq [0,1]$, we have $f\mu(X \cap B) \in U \iff \mu(X) \in U$, and this implies that f is continuous and open. We are left to show bijectivity, but an inverse of f is easily found: define $g \colon \mathfrak{M}(B) \to \mathfrak{M}_x^B(\mathfrak{U})$ as $g\nu(X) = \nu(X \cap B)$, where $g\nu \in \mathfrak{M}_x^B(\mathfrak{U})$ because $S(g\nu) \subseteq B$ by definition. □

As we will see in the following chapter, under NIP a lot of theorems about types generalize to theorem about measures. For the time being, we will just settle on definitions and first results.

Definition 2.103. Let A be a small subset of $\mathfrak{U}$ and $\mu \in \mathfrak{M}_x(\mathfrak{U})$. We call μ

Non-forking over A iff for all $\varphi(x;b) \in L(\mathfrak{U})$ if $\mu(\varphi(x;b)) > 0$ then $\varphi(x;b)$ does not fork over A.

Strongly Lstp_A*-invariant* iff for all $a \equiv_{\mathrm{Lstp}_A} b$ we have $\mu(\varphi(x;a) \triangle \varphi(x;b)) = 0$.

Lstp_A*-invariant* iff for all $a \equiv_{\mathrm{Lstp}_A} b$ we have $\mu(\varphi(x;a)) = \mu(\varphi(x;b))$.

Strongly A-invariant iff for all $a \equiv_A b$ we have $\mu(\varphi(x;a) \triangle \varphi(x;b)) = 0$.

A-invariant iff for all $a \equiv_A b$ we have $\mu(\varphi(x;a)) = \mu(\varphi(x;b))$.

Finitely satisfiable in A iff for all $\varphi(x;b) \in L(\mathfrak{U})$ if $\mu(\varphi(x;b)) > 0$ then $\varphi(x;b)$ has a point in A.

If we do not mention A, it means that the property holds for some small A.

Notice that invariance can be "translated" in two different ways. Needless to say, strong invariance implies invariance, and the distinction does not exist if μ is a type. One of the advantages of the "strongly" definitions is that they correspond to properties of $S(\mu)$:

Proposition 2.104. Let X be one property among "Non-forking over A", "Strongly Lstp_A-invariant", "Strongly A-invariant" and "Finitely satisfiable in A". Let $S_x^X(\mathfrak{U}, A)$ (resp. $\mathfrak{M}_x^X(\mathfrak{U}, A)$) be the closed subspace of types (resp. measures) that are X. Then the following facts hold:

1. μ is X if and only if every $p \in S(\mu)$ is X. In other words, $\mu \in \mathfrak{M}_x^X(\mathfrak{U}, A)$ if and only if $S(\mu) \subseteq S_x^X(\mathfrak{U}, A)$.

2. Each property in the list implies the previous one.

3. $\mathfrak{M}_x^X(\mathfrak{U}, A)$ is homeomorphic to the space of measures[17] on $S_x^X(\mathfrak{U}, A)$.

Proof. All such properties are expressed by saying that μ must give measure zero to certain formulas, and this proves the first statement. The second one follows because we already know that those implications are true at the level of types, and the third one is a special case of Lemma 2.102. □

Generalizing definability is slightly trickier. For the moment, we just give the definition, and we will come back on it later.

Definition 2.105. Given $\mu \in \mathfrak{M}_x^{\mathrm{inv}}(\mathfrak{U}, A)$, $C \subseteq [0,1]$ and $\varphi(x;y)$ in L, let $d_{\mu_x}^C \varphi(x;y) = \{\mathrm{tp}(b/A) \mid \mu(\varphi(x;b)) \in C\}$, which is well-defined by A-invariance. We say that μ is *definable over A* iff for all closed $C \subseteq [0,1]$ and $\varphi(x;y) \in L$, the set $d_\mu^C \varphi$ is closed in $S_y(A)$. We say that μ is *Borel-definable over A* iff for all closed $C \subseteq [0,1]$ and $\varphi(x;y) \in L$, the set $d_\mu^C \varphi$ is Borel in $S_y(A)$.

Extending Measures

When seen as measures on $\mathrm{Def}_x(A)$, Keisler measures are only required to be *finitely* additive. This enables us to extend them to bigger Boolean algebras with great ease:

Proposition 2.106 ([Sim15, Lemma 7.3])**.** Every finitely additive probability measure defined on a sub-Boolean-algebra of $\mathrm{Def}_x(A)$ extends to a Keisler measure over A.

Proof. Let $\mu\colon \Omega \to [0,1]$ be such. Extending it to $\mathrm{Def}_x(A)$ amounts to finding a point in the following closed subset of $\mathfrak{M}_x(A)$

$$\bigcap_{X\in\Omega} \{f \in \mathfrak{M}_x(A) \mid f(X) = \mu(X)\}$$

and by compactness we only have to define an extension f on each finite sub-Boolean-algebra B of $\mathrm{Def}_x(A)$. Write $B = \langle \psi_0, \dots, \psi_{n-1}\rangle$, where the ψ_i are the atoms of B. Clearly it is sufficient to define f on the ψ_i, but we must make sure that f will be compatible with μ. Let $\varphi_0, \dots, \varphi_{m-1}$ be the atoms of the Boolean algebra $\Omega \cap B$. Since this is a sub-Boolean-algebra of B, for each $j < m$ there is a subset $S_j \subseteq n$ such that $\varphi_j = \bigvee_{i \in S_j} \psi_i$. Since $\psi_{i_0} \wedge \psi_{i_1} = \bot$ for all $i_0 \neq i_1$ by definition of atom, all we need to find is a finitely additive f satisfying $\sum_{i\in S_j} f(\psi_i) = \mu(\varphi_j)$, and the only possible obstruction is some ψ_i being an element of S_j for several different $j < m$. Actually this cannot happen: since for each $j_0 \neq j_1$ we have $\varphi_{j_0} \wedge \varphi_{j_1} = \bot$ (because the φ_j are atoms too), then $S_{j_0} \cap S_{j_1} = \emptyset$. □

[17] Recall that with "measure" we mean "regular measure".

While Proposition 2.106 will be a very useful technical result, there is another, more elegant way of finding extensions of measures.

Definition 2.107. Let $\mu \in \mathfrak{M}_x(M)$. Let $[0,1]$ be standard real interval, $<$ the usual relation on it, and $+$ the addition modulo 1. For all $\varphi(x;y) \in L$ define $f_\varphi \colon M \to [0,1]$ as $f_\varphi(b) = \mu(\varphi(x;b))$. Put all these together in a structure

$$\widetilde{M}_\mu = (M, [0,1], <, +, \{f_\varphi \mid \varphi(x;y) \in L\})$$

If $[0,1]^* \succ [0,1]$ an easy analysis of the types over $[0,1]$ yields a standard part map $\mathrm{st}\colon [0,1]^* \to [0,1]$ such that $\mathrm{st}(u)$ is the unique real such that $|u - \mathrm{st}(u)|$ is infinitesimal. It is then easy to check the following fact.

Fact 2.108 ([Sim15, p. 101])**.** For all $\widetilde{N} \succ \widetilde{M}_\mu$ there is an extension $\nu \in \mathfrak{M}_x(\widetilde{N} \restriction L)$ of μ defined by $\nu(\varphi(x;b)) = \mathrm{st}(f_\varphi(b))$.

An even cleaner way of implementing this idea would be to define $\widetilde{M}_\mu$ not as a first-order structure, but as a structure in *continuous logic* (see [BBHU08, BU10,BP10]), in order to have the interval $[0,1]$ built inside the logical apparatus. Continuous logic is also a convenient framework to "realize" measures over a model of a theory T considering the *randomization* of T, which is essentially a theory whose models are made of random variables with values in models of T. See [BK09].

Now that we can extend measures we may ask ourselves how such extensions can behave. For instance, suppose that $\mu \in \mathfrak{M}_x(M)$, $\varphi(x;b) \in L(\mathfrak{U})$, and we have some $\psi(x) \in L(M)$ such that $\vDash \varphi(x;b) \to \psi(x)$. Then, clearly, any ν extending μ must take into account $\mu(\psi(x))$ as an upper bound for $\nu(\varphi(x;b))$. The following result tells us that this is the only constraint, i.e. we can find such a ν with a prescribed value for $\nu(\varphi(x;b))$, as long as this values makes sense. We first give a definition.

Definition 2.109. If B is Borel and $\mu(B) > 0$, the *localization of μ at B* is defined as $\mu_B(\varphi(x)) = \mu([\varphi(x)] \cap B)/\mu(B)$.

Proposition 2.110 ([Sim15, Lemma 7.4])**.** Let $\mu \in \mathfrak{M}_x(M)$ and $\varphi(x;b) \in L(\mathfrak{U})$. Define

$$r_0 = \sup_{\substack{\psi(x) \in L(M) \\ \vDash \psi(x) \to \varphi(x;b)}} \mu(\psi(x)) \qquad\qquad r_1 = \inf_{\substack{\psi(x) \in L(M) \\ \vDash \varphi(x;b) \to \psi(x)}} \mu(\psi(x))$$

Then for all $r \in [r_0, r_1]$ there is $\nu \in \mathfrak{M}_x(\mathfrak{U})$ such that $\nu \restriction M = \mu$ and $\nu(\varphi(x;b)) = r$.

Proof. If we find such a ν_0 for $r = r_0$ and a ν_1 for $r = r_1$, then by convexity we can set $\nu = \frac{r_1 - r}{r_1 - r_0}\nu_1 + \frac{r - r_0}{r_1 - r_0}\nu_2$ and be done. Moreover, up to applying the result to $\neg\varphi(x;b)$, we only have to show it for r_1. By Proposition 2.106

it suffices to define ν_1 on the Boolean algebra Ω generated by $\mathrm{Def}_x(M)$ and $[\varphi(x;b)]$. If $r_1 = 0$, then there is exactly one ν_1 defined on Ω such that $\nu_1 \upharpoonright \mathrm{Def}_x(M) = \mu$ and $\nu_1(\varphi(x;b)) = 0$, and if $r_1 = 1$ we can apply the case $r_1 = 0$ to $\neg\varphi$. Therefore we have proven the proposition in these two special cases. Suppose now that $0 < r_1 < 1$, and consider the closed set

$$C = \bigcap_{\substack{\psi(x) \in L(M) \\ \vDash \varphi(x;b) \to \psi(x)}} [\psi(x)]$$

We can apply to the localization μ_C one of the already proven special cases and get ν_2 such that $\nu_2 \upharpoonright \mathrm{Def}_x(M) = \mu_C$ and $\nu_2(\varphi(x;b)) = 1$. Similarly, we find ν_3 such that $\nu_3 \upharpoonright \mathrm{Def}_x(M) = \mu_{C^\complement}$ and $\nu_3(\varphi(x;b)) = 0$. Setting $\nu = r_1\nu_2 + (1 - r_1)\nu_3$ completes the proof. □

Of course, having unique extensions to all bigger models is a property that deserves a name:

Definition 2.111. We say that $\mu \in \mathfrak{M}_x(M)$ is *smooth* iff for all $N \succ M$ it has a unique extension in $\mathfrak{M}_x(N)$. If $M_0 \prec M$ and $\mu \upharpoonright M_0$ is already smooth, we say that μ is *smooth over* M_0.

Remark 2.112. If we regard a type $p \in S_x(M)$ as a measure, then it is smooth if and only if it is already realized in M: if $a \neq b$ are different realizations of p, then we have one extension containing $x = a$ and another containing $x = b$, but if p has only one realization, then it is algebraic and must be already realized in M.

Proposition 2.110 can be used to characterize smooth measures.

Proposition 2.113 ([Sim15, Lemma 7.8])**.** A measure $\mu \in \mathfrak{M}_x(M)$ is smooth if and only if the following holds. For all $\varphi(x;y) \in L$ and $\varepsilon > 0$ there are $n \in \omega$ and, for $i < n$, some $L(M)$-formulas $\psi_i(y)$, $\theta_i^0(x)$ and $\theta_i^1(x)$ such that

1. $S_y(M) = \bigsqcup_{i<n} [\psi_i(y)]$,
2. for all $i < n$ if $\vDash \psi_i(b)$ then $\vDash \theta_i^0(x) \to \varphi(x;b) \to \theta_i^1(x)$,
3. for all $i < n$ we have $\mu(\theta_i^1(x)) - \mu(\theta_i^0(x)) < \varepsilon$.

Proof. ⇐ Let ν extend μ. We will show that for all $b \in \mathfrak{U}$ and $\varphi(x;y) \in L$ there is only one possible value for $\nu(\varphi(x;b))$. Given $\varepsilon > 0$, apply the hypotheses and find the unique i such that $\vDash \psi(b)$. Then we must have $\nu(\varphi(x;b)) \in (\mu(\theta_i^0(x)), \mu(\theta_i^1(x)))$, and this interval has length smaller than ε. Clearly the intersection of such intervals as $\varepsilon \to 0$ contains exactly one point.

(⇒) For a fixed b, the existence of some θ^0 and θ^1 that work is due to Proposition 2.110, since if there were none it would mean that we have an interval of length at least ε where to choose values for extending of μ to $\varphi(x;b)$. Then $\operatorname{tp}_y(b/M)$ proves $\forall x\ \theta^0(x) \to \varphi(x;y) \to \theta^1(x)$, and by compactness there is $\psi(y) \in \operatorname{tp}_y(b/M)$ already implying it. Using compactness again, we can cover $S_y(M)$ with finitely many ψ_i, and up to replacing $[\psi_i]$ with $\left[\psi_i \wedge \neg \bigvee_{j\neq i} \psi_j\right]$ and discarding the empty ones we get a partition. □

Approximation of Smooth Measures

Using the previous characterization it is possible to approximate smooth measures with averages of types.

Notation 2.114. We write $\operatorname{Av}(\mu_0,\dots,\mu_{n-1};B)$ for $\frac{1}{n}\sum_{k<n}\mu_k(B)$. Similarly $\operatorname{Av}(a_0,\dots,a_{n-1};B)$ denotes $|\{i<n \mid\models a_i \in B\}|/n$. In other words, it is $\operatorname{Av}(\mu_0,\dots,\mu_{n-1};B)$ where $\mu_i = \operatorname{tp}(a_i/\mathfrak{U})$.

Proposition 2.115 ([HPS13, Corollary 2.6])**.** Let $\mu \in \mathfrak{M}_x(\mathfrak{U})$ be smooth over a small M. For all $B_0,\dots,B_{m-1}$ Borel subsets of $S_x(M)$, $\varphi(x;y) \in L$ and $\varepsilon > 0$ there are $a_0,\dots,a_{n-1} \in \mathfrak{U}$ such that for all $i < n$ we have $\operatorname{tp}(a_i/\mathfrak{U}) \in S(\mu)$ and for all $b \in \mathfrak{U}$ and $j < m$

$$|\mu(B_j \cap \varphi(x;b)) - \operatorname{Av}(a_0,\dots,a_{n-1};B_j \cap \varphi(x;b))| \le \varepsilon$$

Proof. Proposition 2.113 applied to $\mu \restriction M$, $\varphi(x;y)$ and $\varepsilon/8$ gives us $\ell \in \omega$ and, for $h < \ell$, some $\psi_h(y), \theta_h^0(x), \theta_h^1(x) \in L(M)$ that we can use to approximate $\mu(\varphi(x;b))$.

Claim. There are some n and, for $i < n$, some $p_i \in S(\mu \restriction M) \subseteq S_x(M)$ such that for all $h < \ell$, $j < m$ and $k \in \{0,1\}$

$$|\mu(B_j \cap \theta_h^k(x)) - \operatorname{Av}(p_0,\dots,p_{n-1};B_j \cap \theta_h^k(x))| \le \frac{\varepsilon}{4}.$$

Proof of the Claim. By Proposition A.35 we can find some probability space (Ω,Σ,P) and a sequence of independent measurable functions $\{Y_i\colon \Omega \to S(\mu \restriction M) \mid i < \omega\}$, all with law μ, i.e. such that $(Y_i)^*(P) = \mu$. Letting $\chi(B)$ denote the characteristic function of B, define

$$\tilde{Y}_i^{h,j,k} = \chi\big(B_j \cap \theta_h^k(x)\big) \circ Y_i \colon \Omega \to [0,1] \qquad Z_N^{h,j,k} = \sum_{i<N} \tilde{Y}_i^{h,j,k}$$

Notice that, denoting with $\mathbb{E}$ the expected value and using Lemma 2.101,

$$\mathbb{E}(Z_N^{h,j,k}) = N\cdot\mathbb{E}(\tilde{Y}_0^{h,j,k}) = N\cdot\mu(B_j\cap\theta_h^k(x)\cap S(\mu\restriction M)) = N\cdot\mu(B_j\cap\theta_h^k(x))$$

and in particular $\tilde{Y}_i^{h,j,k} \in L^1$. By a version of the Law of the Large Numbers, namely Theorem A.36, for all h, j, k the sequence $(Z_N^{h,j,k} - \mathbb{E}(Z_N^{h,j,k}))/N$ converges in probability to 0. Hence, if we denote

$$A_N^{h,j,k} = \{\rho \in \Omega \mid Z_N^{h,j,k}(\rho)/N - \mu(B_j \cap \theta_h^k(x)) > \varepsilon/4\}$$

then there is $N_{h,j,k}$ such that for all $n > N_{h,j,k}$ we have $P(A_n^{h,j,k}) < 1/2m\ell$. Let $n = \max_{h,j,k} N_{h,j,k}$ and $A = \bigcup_{h,j,k} A_n^{h,j,k}$. Then

$$P(A) \leq \sum_{\substack{h<\ell \\ j<m \\ k<2}} P(A_n^{h,j,k}) \leq \sum_{\substack{h<\ell \\ j<m \\ k<2}} P(A_{N_{h,j,k}}^{h,j,k}) < \sum_{\substack{h<\ell \\ j<m \\ k<2}} \frac{1}{2m\ell} = 1$$

Therefore $A^{\complement}$ has positive probability, hence it is non-empty. Let $\rho \in A^{\complement}$ and, for $i < n$, set $p_i = Y_i(\rho)$; then, for all h, j, k, we have

$$\frac{Z_N^{h,j,k}(\rho)}{N} = \frac{1}{N} \sum_{i<N} \chi\big(B_j \cap \theta_h^k(x)\big) \circ Y_i(\rho) = \operatorname{Av}(p_0, \ldots, p_{n-1}; B_j \cap \theta_h^k(x))$$

which is no further than $\varepsilon/4$ from $\mu(B_j \cap \theta_h^k(x))$ by construction. □ CLAIM

Let the $p_i \in S_x(M)$ be given by the Claim and, since M is small, realize them with suitable $a_i \in \mathfrak{U}$, which we may suppose without loss of generality[18] to be such that $\operatorname{tp}(a_i/\mathfrak{U}) \in S(\mu)$; then set $\nu(B) = \operatorname{Av}(a_0, \ldots, a_{n-1}; B)$. Now let $b \in \mathfrak{U}$ and $j < m$, and let $h < \ell$ be the unique one such that $b \vDash \psi_h(y)$. By choice of h, for all $j < m$ we have

$$|\mu(B_j \cap \varphi(x; b)) - \mu(B_j \cap \theta_h^0(x))| \leq |\mu(\theta_h^1(x)) - \mu(\theta_h^0(x))| \leq \frac{\varepsilon}{8}$$

Hence, using the Claim and denoting $\varphi(x; b)$ and $\theta_h^k(x)$ with φ and θ_h^k,

$$\begin{aligned} |\nu(B_j \cap \varphi) - \nu(B_j \cap \theta_h^0)| &\leq |\nu(B_j \cap \theta_h^1) - \nu(B_j \cap \theta_h^0)| \\ &\leq |\nu(B_j \cap \theta_h^1) - \mu(B_j \cap \theta_h^1)| + |\mu(B_j \cap \theta_h^1) - \mu(B_j \cap \theta_h^0)| \\ &\quad + |\mu(B_j \cap \theta_h^0) - \nu(B_j \cap \theta_h^0)| \leq \frac{\varepsilon}{4} + \frac{\varepsilon}{8} + \frac{\varepsilon}{4} = \frac{5}{8}\varepsilon \end{aligned}$$

Therefore, using the Claim once again,

$$\begin{aligned} |\nu(B_j \cap \varphi) - \mu(B_j \cap \varphi)| &\leq |\nu(B_j \cap \varphi) - \nu(B_j \cap \theta_h^0)| + |\nu(B_j \cap \theta_h^0) - \mu(B_j \cap \theta_h^0)| \\ &\quad + |\mu(B_j \cap \theta_h^0) - \mu(B_j \cap \varphi)| \leq \frac{5}{8}\varepsilon + \frac{\varepsilon}{4} + \frac{\varepsilon}{8} = \varepsilon \quad \square \end{aligned}$$

The name "smooth" comes from the fact that one can define a notion of *border* of a formula and smooth measures will be exactly the ones that give measure zero to all borders. See [ABss].

[18] Since $p_i \in S(\mu \restriction M)$, if we regard it as a closed subset of $S_x(\mathfrak{U})$ it meets $S(\mu)$.

Measures and Forking

We conclude by showing that non-dividing can be witnessed by a measure, a fact that will prove itself to be very useful.

Proposition 2.116 ([Sim15, Lemma 7.5])**.** Let $(b_i)_{i<\omega} \in M$ be indiscernible, $\varphi(x;y) \in L$ and $\mu \in \mathfrak{M}_x(M)$. If there is $r > 0$ such that $\mu(\varphi(x;b_i)) \geq r$ for all $i \in \omega$, then $\{\varphi(x;b_i) \mid i < \omega\}$ is consistent.

Proof. Consider $(b_i)_{i<\omega}$ inside $\widetilde{M_\mu}$. Applying the Standard Lemma we can find some $L(\widetilde{M_\mu})$-indiscernible $(c_i)_{i<\omega}$ that still satisfies the hypotheses and extends the Ehrenfeucht-Mostowski type of $(b_i)_{i<\omega}$ over $L(M)$, therefore it is enough to prove the consistency of $\pi(x) = \{\varphi(x;c_i) \mid i < \omega\}$. By $L(\widetilde{M_\mu})$-indiscernibility if $i_0 < \ldots < i_{n-1}$ and $j_0 < \ldots < j_{n-1}$ we have

$$\mu\Big(\bigwedge_{k<n} \varphi(x;c_{i_k})\Big) = \mu\Big(\bigwedge_{k<n} \varphi(x;c_{j_k})\Big)$$

If π is ℓ-inconsistent, then for $n = \ell$ the quantity above is zero. Let ℓ_0 be the maximum n such that $\pi(x)$ is not n-inconsistent, and for $m \in \omega$ define $\psi_m(x) = \bigwedge_{k<\ell_0} \varphi(x;c_{m\ell_0+k})$. Since $\mu(\varphi(x;c_i)) \geq r > 0$ by hypothesis, $\ell_0 > 0$, and thus for $m_0 \neq m_1$ we have $\psi_{m_0} \neq \psi_{m_1}$. Then maximality of ℓ_0 implies that $\mu(\psi_{m_0}(x) \wedge \psi_{m_1}(x)) = 0$. This yields $\mu\Big(\bigvee_{m<C} \psi_m\Big) = C\mu(\psi_0)$ and, since by construction $\mu(\psi_m) > 0$, taking $C > 1/\mu(\psi_0)$ contradicts $\mu(x = x) = 1$. $\square$

Corollary 2.117. If A is small and $\mu \in \mathfrak{M}_x(\mathfrak{U})$ is Lstp_A-invariant, then μ is A-non-forking.

Proof. Since A is small it is sufficient to prove that if $\mu(\varphi(x;b)) > 0$ then $\varphi(x;b)$ does not divide over A. Let $(b_i)_{i<\omega}$ be A-indiscernible. By Lstp_A-invariance, for all $i < \omega$ we have $\mu(\varphi(x;b)) = \mu(\varphi(x;b_i))$, and by Proposition 2.116 $\{\varphi(x;b_i) \mid i < \omega\}$ is consistent. $\square$

Chapter 3

NIP Theories

We are now going to introduce the NIP hypothesis and see how it influences invariant types and measures. Thereafter we study *honest definitions*, a powerful tool that will allow us to prove that, if M is NIP, then M^{ext} eliminates quantifiers and inherits dependency from M. We also show how the NIP guarantees the existence of G^{00} (see Definition 1.81) and its preservation when passing to M^{ext}, and conclude with the existence of invariant heirs. The references are mainly [Sim15, CS13, HPP08, CPS14].

3.1 The Independence Property

This section is devoted to definitions and fundamental results.

VC-Dimension

Definition 3.1. A family $\mathcal{S} \subseteq \mathscr{P}(X)$ of parts of X *shatters* a subset A of X iff $\mathscr{P}(A) = \{A \cap S \mid S \in \mathcal{S}\}$.

In other words, $\mathcal{S}$ shatters A if the whole power-set of A can be obtained "looking at A through members of $\mathcal{S}$".

Example 3.2 ([Sim15, Example 2.4])**.** Any subset A of a totally ordered set O shattered by the family of closed right half-lines $\{x \geq b \mid b \in O\}$ has at most one point: if $a_0 < a_1$ are in A, then $\{a_0\}$ cannot be written as the intersection of A with a right half-line.

Example 3.3. The power-set of X trivially shatters any $A \subseteq X$.

Example 3.4. The family of all bounded regions delimited by a Jordan curve in $\mathbb{R}^2$ shatters any finite subset of $\mathbb{R}^2$.

Example 3.5. Let $\mathcal{S}$ be the family of rectangles in $\mathbb{R}^2$ whose sides are parallel to the axes, i.e. the family of the $S \subseteq \mathbb{R}^2$ of the form

$$(x_0, x_1) \in S \iff a_0 \leq x_0 \leq b_0 \wedge a_1 \leq x_1 \leq b_1$$

Figure 3.1: A Jordan curve and a rectangle isolating subsets of points.

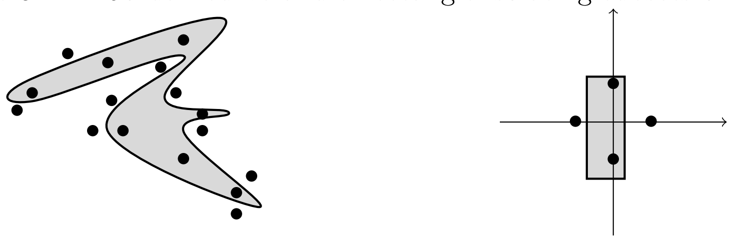

where $a_i < b_i \in \mathbb{R}$. Then $\mathcal{S}$ shatters the set of vertexes of the $\|\cdot\|_1$-ball, i.e. $\{(1,0),(0,1),(-1,0),(0,-1)\}$.

Definition 3.6. The VC-*dimension* of a family $\mathcal{S} \subseteq \mathscr{P}(X)$ is either the maximum $n \in \omega$ such that $\mathcal{S}$ shatters a subset of X of cardinality n, or ∞ if $\mathcal{S}$ shatters finite sets of arbitrarily large cardinality.

Remark 3.7. If $\mathcal{S}$ shatters a set A of cardinality n, then it automatically shatters all the subsets of A, hence it shatters a set of cardinality m for all $m \leq n$.

Example 3.8. The family of Example 3.5 has VC-dimension 4. We have already exhibited a 4-element set shattered by $\mathcal{S}$, and by the previous remark it suffices to show that $\mathcal{S}$ shatters no 5-element one. Let $|A| = 5$ and fix a 4-element subset B of A consisting of one point minimizing x_0, one maximizing it, one minimizing x_1 and one maximizing it; in other words the first one has no points of A strictly to the left, the second one has none strictly to the right, etc. This set cannot be written as $A \cap S$ for any rectangle S with horizontal sides, because $A \cap S$ will inevitably contain the point in $A \setminus B$ or miss some point of B.

Dependent Formulas

A formula can be thought as a uniformly definable family of definable subsets of a model as follows. We partition its free variables in "set variables", which are intended to range over points of the definable set, and "parameter variables", which are intended to code members of the family. For instance we can describe the family $\mathcal{S}$ of Example 3.5 with the partitioned formula

$$\varphi(x;y) = \varphi(x_0, x_1; y_0, y_1, y_2, y_3) = (y_0 \leq x_0 \leq y_1) \wedge (y_2 \leq x_1 \leq y_3)$$

In other words, we can associate to $\varphi(x;y)$ the family $\Phi = \{\varphi(x;b) \mid b \in \mathfrak{U}\}$, and investigate its VC-dimension. The NIP formulas will be[1] the ones where

[1] The compactness argument is easy, but for the sake of clarity and at the risk of proving the obvious we will take Definition 3.9 as the "official" definition of NIP formula and prove the equivalence mentioned in this paragraph in Lemma 3.12.

this VC-dimension is finite. Since, by compactness and saturation, Φ shatters arbitrarily large sets if and only if it shatters and infinite subset of $\mathfrak{U}$, we give the definition this way:

Definition 3.9. The *partitioned* formula $\varphi(x; y)$ has the *independence property*, or IP, iff there are two infinite sets $A = \{a_i \mid i \in I\} \subseteq \mathfrak{U}^{|x|}$ and $B = \{b_X \mid X \in \mathscr{P}(I)\} \subseteq \mathfrak{U}^{|y|}$ such that $\vDash \varphi(a_i; b_X) \iff i \in X$. A formula has NIP iff it does not have the independence property.

Obviously enough, NIP stands for No Independence Property/Not the Independence Property. We also say that the formula *is* NIP, or *dependent*.

Remark 3.10. The word "partitioned" is highlighted in the previous definition because this *does* depend on where the semicolon is. For instance, if $\varphi(x; y)$ has IP, and we re-baptize it $\psi(x, y, z) = \varphi(x, y)$, then $\psi(x; y, z)$ has IP, but $\psi(x, y; z)$ does not.

Lemma 3.11. If $\{\varphi(x; b) \mid b \in \mathfrak{U}\} \subseteq \mathscr{P}(\mathfrak{U})$ has infinite VC-dimension and I is an infinite set with small size[2] with respect to $\mathfrak{U}$, then there are two infinite sets $A = \{a_i \mid i \in I\} \subseteq \mathfrak{U}^{|x|}$ and $B = \{b_X \mid X \in \mathscr{P}(I)\} \subseteq \mathfrak{U}^{|y|}$ such that $\vDash \varphi(a_i; b_X) \iff i \in X$.

Proof. Saying that $\{\varphi(x; b) \mid b \in \mathfrak{U}\}$ has infinite VC-dimension means that, for all $n \in \omega$, there is an n-element set $A_n = \{\tilde{a}_i^n \mid i < n\}$ and some $B_n = \{\tilde{b}_X^n \mid X \in \mathscr{P}(n)\}$ such that $\mathfrak{U} \vDash \varphi(\tilde{a}_i^n; \tilde{b}_X^n)$ if and only if $i \in X$. Take a small $M \prec \mathfrak{U}$ containing $\bigcup_{n\in\omega} A_n \cup B_n$, then let $\{c_i \mid i \in I\}$ and $\{d_X \mid X \in \mathscr{P}(I)\}$ be two new sets of constants. Then the union of the elementary diagram of M with the theory below is finitely satisfiable:

$$\{\varphi(c_i; d_X) \mid i \in I, X \in \mathscr{P}(I), i \in X\} \cup \{\neg\varphi(c_i; d_X) \mid i \in I, X \in \mathscr{P}(I), i \notin X\}$$

By compactness it has a model, and by Löwenheim-Skolem and the fact that I is small we can suppose it has small cardinality, so we can embed its $L(M)$-reduct into $\mathfrak{U}$. Then taking as a_i and b_X the interpretations of c_i and d_X ends the proof. □

Lemma 3.12. $\varphi(x; y)$ has NIP if and only if $\{\varphi(x; b) \mid b \in \mathfrak{U}\} \subseteq \mathscr{P}(\mathfrak{U})$ has finite VC-dimension.

Proof. (⇒) Apply Lemma 3.11 with your favourite infinite small I.

(⇐) If $\varphi(x; y)$ had IP as witnessed by $A = \{a_i \mid i \in I\}$ and $B = \{b_X \mid X \in \mathscr{P}(I)\}$ by some infinite I, then $\{\varphi(x; b) \mid b \in \mathfrak{U}\}$ would shatter the infinite set A, and hence all of its finite subsets, thus having infinite VC-dimension. □

[2]Recall (see Section A.6) that this also means that $\mathscr{P}(I)$ has small size.

Notation 3.13. The *opposite formula* of $\varphi(x;y)$ is the same formula as φ, but the y are now considered the "set variables" and the x the "parameter" ones. Since we usually write set variables on the left of the semicolon and parameter variables on the right, we will use the notation $\varphi^{\mathrm{opp}}(y;x) = \varphi(x;y)$.

Despite Remark 3.10, being NIP does not depend on which side of the semicolon variables are:

Proposition 3.14 ([Sim15, Lemma 2.5])**.** A partitioned formula is NIP if and only if its opposite formula is.

Proof. If $\varphi(x;y)$ has IP, by Lemma 3.12 and Lemma 3.11 we can find inside $\mathfrak{U}$ two subsets $A = \{a_X \mid X \in \mathscr{P}(\omega)\}$ and $B = \{b_{\mathcal{F}} \mid \mathcal{F} \in \mathscr{P}(\mathscr{P}(\omega))\}$ such that $\mathfrak{U} \vDash \varphi(a_X; b_{\mathcal{F}})$ if and only if $X \in \mathcal{F}$. Let $\sqcup_i$ be the principal ultrafilter on i, i.e. $\{X \in \mathscr{P}(\omega) \mid i \in X\}$. Then $\varphi^{\mathrm{opp}}(y;x)$ shatters the infinite set $\{\sqcup_i \mid i \in \omega\} \subseteq B$, because for all $i \in \omega$ and $X \in \mathscr{P}(\omega)$

$$i \in X \iff X \in \sqcup_i \iff \vDash \varphi(a_X; b_{\sqcup_i}) \iff \vDash \varphi^{\mathrm{opp}}(b_{\sqcup_i}; a_X) \qquad \square$$

The VC-dimension of φ^{opp} need not equal the one of φ, but some bounds can be proven; see for instance [Sim15, Lemma 6.3]. For the purposes of this thesis, the only relevant bound will be $< \infty$, and we just said that this holds for φ if and only if it holds for φ^{opp}.

Permanent Assumption 3.15. Henceforth Lemma 3.11, Lemma 3.12 and Proposition 3.14 will be used even without explicit mention.

Let us see some examples of formulas with or without NIP.

Example 3.16. We saw that the family of Example 3.5 has VC-dimension 4, hence any formula defining it has NIP.

Example 3.17 ([Sim15, Example 2.4])**.** In the random graph, the formula $\varphi(x;y) = E(x,y)$ shatters arbitrarily large finite sets, as it follows easily from the random graph axioms.

Example 3.18 ([Sim15, Example 2.4])**.** In $\mathrm{Th}(\mathbb{N})$, the formula $\varphi(x;y) = x \mid y$, i.e. "x divides y", shatters any finite set of prime numbers $\{p_i \mid i < n\}$: to isolate $\{p_i \mid i \in X\}$ just take as y the product $\prod_{i \in X} p_i$.

Example 3.19 ([New09, Example 1])**.** Let $\eta \in 2^{\mathbb{Z}}$ be a concatenation of all finite strings over $\{0,1\}$ as in Example 1.12. Expand the structure $(\mathbb{Z}, +)$ adding a predicate $P_\sigma^{(1)}$ for all finite $\{0,1\}$-strings σ, true for n iff[3] the string $\eta(0-n), \eta(1-n), \dots, \eta(|\sigma|-1-n)$ coincides with σ. Then if σ is the string "1", the formula $\varphi(x;y) = P_\sigma(x+y)$ has IP.

[3] In other words, $P_\sigma(n)$ is true if, after n right shifts s of η, we look at $s^{(n)}\eta(0), \dots, s^{(n)}\eta(|\sigma|-1)$ and read σ.

For instance, we can shatter a 3-element set in the following way: find a_0, a_1, a_2 such that shifting by them produces the following configurations, where the y-th column[4] carries at row i the value of $\eta(a_i + y)$

$$\begin{array}{lcccccccccc} a_0: & \dots & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & \dots \\ a_1: & \dots & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 1 & \dots \\ a_2: & \dots & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & \dots \end{array}$$

and in order to select, say, $\{a_0, a_2\}$, one considers $P_\sigma(x+5)$. The matrix above basically codes $\mathscr{P}(3)$ and should hint that this trick can be generalized to shatter an arbitrarily large finite set.

The following characterization will be of fundamental importance, and actually some authors take it as the definition of NIP.

Theorem 3.20 ([Sim15, Lemma 2.7])**.** A partitioned formula $\varphi(x;y) \in L(A)$ has IP if and only if there are an A-indiscernible sequence $(a_i)_{i<\omega}$ and a b in $\mathfrak{U}$ such that $\varphi(a_i;b)$ is true for all and only the even i, i.e.

$$\mathfrak{U} \vDash \varphi(a_i;b) \iff \exists j \in \omega \; i = 2j$$

Proof. ($\Rightarrow$) Let $(\tilde{a}_i)_{i<\omega}$ be shattered by φ, and by the Standard Lemma let $(a_i)_{i<\omega}$ be an A-indiscernible sequence realizing its Ehrenfeucht-Mostowski type over A. Then the partial type $\pi(x) = \{\varphi(a_{2i};y) \mid i \in \omega\} \cup \{\neg\varphi(a_{2i+1};y) \mid i \in \omega\}$ is consistent, since for all $n \in \omega$ the formula $\exists y \bigwedge_{i<n} \varphi(x_{2i};y) \wedge \neg\varphi(x_{2i+1};y)$ is in $\mathrm{EM}(\tilde{a}_i)_{i<\omega} \subseteq \mathrm{EM}(a_i)_{i<\omega}$. Any $b \vDash \pi(x)$ satisfies the thesis.

($\Leftarrow$) Let $X \in \mathscr{P}(\omega)$ be infinite, and fix an increasing bijection τ from X to the evens. Since τ is increasing, by A-indiscernibility $\{a_i \mapsto a_{\tau(i)} \mid i < \omega\}$ extends to an $f \in \mathrm{Aut}(\mathfrak{U}/A)$. Then we have

$$i \in X \iff \vDash \varphi(f(a_i);b) \iff \vDash \varphi(f^{-1}(f(a_i)); f^{-1}(b)) \iff \vDash \varphi(a_i; f^{-1}(b))$$

Now fix $n \in \omega$. Since the family of infinite subsets of ω shatters n, we can use the previous construction to shatter a suitable n-element subset of $\{a_i \mid i < \omega\}$. □

Definition 3.21. Let $\varphi(x;y)$ be a partitioned $L(A)$-formula, $(a_i)_{i\in I}$ an A-indiscernible sequence and $b \in \mathfrak{U}$. The *number of alternations of* $\varphi(x;b)$ *on* $(a_i)_{i\in I}$ is either the maximum n such that there are $i_0 < \dots < i_n$ such that for all $k < n$ we have $\vDash \varphi(a_{i_k}) \triangle \varphi(a_{i_{k+1}})$, or ∞ if there are arbitrarily large such n; we denote it with $\mathrm{alt}(\varphi(x;b), (a_i)_{i\in I})$. The *alternation rank* of $\varphi(x;y)$ is

$$\mathrm{alt}(\varphi(x;y)) = \sup_{\substack{b\in\mathfrak{U} \\ (a_i)_{i\in I}\ A\text{-indiscernible}, a_i\in\mathfrak{U}}} \mathrm{alt}(\varphi(x;b), (a_i)_{i\in I})$$

[4]The leftmost column corresponds to $y = 0$.

Corollary 3.22 ([Sim15, Proposition 2.8])**.** The formula $\varphi(x;y) \in L(A)$ is NIP if and only if $\mathrm{alt}(\varphi(x;y))$ is finite, if and only if for all A-indiscernible $(a_i)_{i\in I}$ and $b \in \mathfrak{U}$ there is an end segment of I where $\varphi(a_i;b)$ is either always true or always false.

Proof. Immediate from Theorem 3.20. □

Permanent Assumption 3.23. Henceforth we add Theorem 3.20 and Corollary 3.22 to the list of results that will be used without explicit mention.

Lemma 3.24 ([Sim15, Lemma 2.9])**.** In all theories, NIP formulas are closed under Boolean combinations.

Proof. Closure under negation is obvious. If both $\varphi(x;y)$ and $\psi(x;y)$ are NIP $L(A)$-formulas, take any A-indiscernible $(a_i)_{i\in I}$ and any $b \in \mathfrak{U}$. Then for some $i_\varphi \in I$, for all $i \geq i_\varphi$ the truth value of $\varphi(a_i;b)$ is constant, and so is, for some $i_\psi \in I$, the truth value of $\psi(a_i;b)$ for all $i \geq i_\psi$. Then the same happens with $\varphi \wedge \psi$ setting $i_{\varphi\wedge\psi} = \max\{i_\varphi, i_\psi\}$. □

Dependent Theories

Definition 3.25. A theory is NIP iff all of its formulas are NIP. As structure is NIP iff its theory is.

Example 3.26. Stable theories are NIP.

Proof. One of the equivalent definitions of instability (see for instance [Sim15, Proposition 2.55]) is the existence of some $\varphi(x;y)$, some $(a_i)_{i<\omega}$ and some $(b_j)_{j<\omega}$ such that $\vDash \varphi(a_i;b_j) \iff i < y$. If $\varphi(x;y)$ has IP as witnessed by $(a_i)_{i<\omega}$ and $(b_X)_{X\in\mathscr{P}(\omega)}$, it suffices to set $b_j = b_{\{i\in\omega \mid i<j\}}$ to witness instability. □

Remark 3.27. Outside some special cases such as $I = \omega$, Corollary 3.22 does *not* imply that $\{i \in I \mid \vDash \varphi(a_i;b)\}$ is either finite or cofinite, but it *does* imply that it is a finite union of intervals[5]. For instance, we could have $I = \mathbb{R}$ and $\{i \in I \mid \vDash \varphi(a_i;b)\} = (-\infty,\pi) \cup \{5\} \cup [7,8)$. Indeed, $\{i \in I \mid \vDash \varphi(a_i;b)\}$ always being finite or cofinite is equivalent to stability. This can be easily shown exploiting Corollary 3.22 and the fact that[6] a theory is stable if and only if every permutation of an A-indiscernible sequence is still A-indiscernible.

Lemma 3.28. The NIP is preserved under adding parameters and taking reducts. In other words, if M is a model of a NIP theory and $A \subseteq M$, then also $\mathrm{Th}_{L(A)}(M)$ is NIP, and if $L_0 \subseteq L$ then $\mathrm{Th}_{L_0}(M)$ is NIP.

[5] Possibly of length 0, i.e. points count as intervals.

[6] It follows easily from the definition of stability mentioned in the proof of Example 3.26, anyway see [Sim15, Lemma 2.59].

Proof. If $\varphi(x;y) = \psi(x,y,c_a)$ is an $L(A)$-formula in $\mathrm{Th}_{L(A)}(M)$ with IP, where $\psi(x,y,z) \in L$ and c_a is a tuple of constant symbols corresponding to the tuple $a \in A$, then $\psi(x;y,z)$ has IP in $\mathrm{Th}_L(M)$. The second statement is trivial. □

We are now going to prove that whether a theory is NIP or not is already decided in dimension one. We first need a lemma.

Lemma 3.29. Suppose that all formulas $\varphi(x;y)$ with $|x| = 1$ are NIP. Then for all $\emptyset$-indiscernible $(a_i \mid i < |T|^+)$ and $b \in \mathfrak{U}^1$ there is $\alpha < |T|^+$ such that $(a_i \mid \alpha < i < |T|^+)$ is $\{b\}$-indiscernible.

Proof. If not, then for all $\alpha < |T|^+$ there are $\delta_\alpha(x_0, \dots, x_{k(\alpha)-1}; y)$ and two increasing sequences of indexes $\alpha < \bar{\imath}(\alpha) < |T|^+$ and $\alpha < \bar{\jmath}(\alpha) < |T|^+$ with $|\bar{\imath}(\alpha)| = |\bar{\jmath}(\alpha)| = k(\alpha)$ such that $\vDash \delta_\alpha(a_{\bar{\imath}(\alpha)}; b) \wedge \neg\delta_\alpha(a_{\bar{\jmath}(\alpha)}; b)$. By pigeonhole, since successor cardinals are regular, up to discarding co-$|T|^+$-many α we may assume that $\delta(\alpha)$ and $k(\alpha)$ do not depend on α and call them δ and k. This allows us to find recursively, for $h < \omega$, some $\bar{\ell}_h$ with $|\bar{\ell}_h| = k$ such that $\delta(a_{\bar{\ell}_h}; b)$ holds if and only if h is even, against the fact that $\delta^{\mathrm{opp}}(y; \bar{x})$ should be NIP because $|y| = 1$. □

Proposition 3.30 ([Sim15, Proposition 2.11])**.** If all formulas in T of the form $\varphi(x;y)$ with $|x| = 1$ are NIP, then T is NIP.

Proof. Fix $\varphi(x;y) = \varphi(x_0, \dots, x_{m-1}; y_0, \dots, y_{n-1})$, an $\emptyset$-indiscernible sequence $(a_i)_{i<|T|^+}$, and $b = (b_0, \dots, b_{n-1})$. By the previous lemma there is α_0 such that cutting away the first α_0 elements of $(a_i)_{i<|T|^+}$ returns a $\{b_0\}$-indiscernible sequence. This is equivalent to the fact that $((a_i, b_0) \mid \alpha_0 < i < |T|^+)$ is $\emptyset$-indiscernible. Applying this reasoning another $n-1$ times returns an α such that $((a_i, b) \mid \alpha < i < |T|^+)$ is $\emptyset$-indiscernible, and therefore the truth value of $\varphi(a_i; b)$ is constant for $i > \alpha$.

This shows that no indiscernible sequence indexed on $|T|^+$ can contradict NIP, so we are left to show the same thing replacing $|T|^+$ with an arbitrary total order I. Up to reversing I and/or eliminating maxima, we may suppose that I has infinite cofinality. We can reduce to the previous case either by compactness, if $\mathrm{cof}(I) < |T|^+$, or by extracting a suitable sub-sequence of length $|T|^+$ if $\mathrm{cof}(I) \geq |T|^+$. □

This is very useful when investigating the dependency of a theory. Here is an example.

Corollary 3.31 ([Sim15, Example 2.12])**.** O-minimal theories are NIP.

Proof. By Proposition 3.30 and Lemma 3.24 we only need to check families of closed right half-lines, and we already did in Example 3.2. □

Theorem 3.32. If M is NIP, then so is M^{eq}.

Proof (Adapted from [TZ12, Lemma 8.4.8 (3)]). Let $\varphi(x; y)$ be a formula in T^{eq} with the IP, with x of sort $\mathfrak{U}^{|x|}/E_x$ and y of sort $\mathfrak{U}^{|y|}/E_y$, for suitable equivalence relations E_x and E_y. Then there are $\{a_i \mid i \in \omega\}$ and $\{b_X \mid X \in \mathscr{P}(\omega)\}$ from $\mathfrak{U}$ such that $\vDash \varphi(\pi_{E_x}(a_i); \pi_{E_y}(b_X)) \iff i \in X$. By Proposition A.51 the L^{eq}-formula $\varphi(\pi_{E_x}(\tilde{x}); \pi_{E_y}(\tilde{y}))$ is equivalent to some L-formula $\psi(\tilde{x}; \tilde{y})$, that will hence have IP. □

3.2 Invariant Types

Borel-Definability

Even if invariant types need not be definable, i.e. $d_p\varphi$ need not be clopen, we will now see that, in NIP theories, it cannot be too complicated either.

Lemma 3.33 ([HP11, Claim in Proposition 2.6])**.** Let $p \in S^{\mathrm{inv}}(\mathfrak{U}, A)$, $b \in \mathfrak{U}$ and suppose that $\varphi(x; y) \in L$ is NIP. Then $\varphi(x; b) \in p$ if and only if there is $n \le \mathrm{alt}(\varphi(x; y))$ such that the following two facts hold.

1. There is $(a_i)_{i<n} \vDash p^{(n)} \restriction A$ such that $\vDash \bigwedge_{i<n-1} \varphi(a_i; b) \triangle \varphi(a_{i+1}; b)$ and $\vDash \varphi(a_{n-1}; b)$.

2. There is no $(\tilde{a}_i)_{i<n+1} \vDash p^{(n+1)} \restriction A$ such that $\vDash \bigwedge_{i<n} \varphi(\tilde{a}_i; b) \triangle \varphi(\tilde{a}_{i+1}; b)$.

Proof. (⇒) Let $(c_j \mid j < \omega) \vDash p^{(\omega)} \restriction A$ realize the maximum number of alternations of truth values for $\varphi(x; b)$ among Morley sequences of p over A. Call this number n and notice that it is at most $\mathrm{alt}(\varphi(x; y))$. By maximality of n, point 2 holds. To show that point 1 holds, let $j_0 < \ldots < j_{n-1}$ be such that $\vDash \varphi(c_{j_i}; b) \triangle \varphi(c_{j_i+1}; b)$ and set $a_i = c_{j_i}$, so we only have to show that $\vDash \varphi(c_{j_{n-1}}; b)$. Let $c_\omega \vDash p \restriction A \cup \{c_i \mid i < \omega\} \cup \{b\}$. Since $\varphi(x; b) \in p$, then by Corollary 2.26 $\vDash \varphi(c_\omega; b)$; hence $\vDash \neg\varphi(c_{j_{n-1}}; b)$ would contradict maximality of n.

(⇐) If $\varphi(x; b)$ is not in p, then $\neg\varphi(x; b)$ is, so we can apply the previous implication to $\neg\varphi(x; b)$ and obtain some $n_{\neg\varphi}$ that satisfies 1 and 2 with respect to $\neg\varphi(x; b)$. Then for any n_φ, either o $n_\varphi < n_{\neg\varphi}$ and thus point 2 cannot hold for n_φ, or $n_\varphi \ge n_{\neg\varphi}$. In this case, since $\neg\varphi(x; b) \in p$, any realization of $p^{(n+1)} \restriction A \cup \{a_i \mid i < n\} \cup \{b\}$ shows that, again, point 2 cannot hold for n_φ. □

Theorem 3.34 ([Sim15, Lemma 7.18], T NIP)**.** Every $p \in S^{\mathrm{inv}}(\mathfrak{U}, A)$ is *strongly Borel-definable over* A, i.e. every $d_{p(x)}\varphi(x; y)$ is a constructible[7] subset of $S_y(A)$.

[7] I.e. a finite Boolean combination of closed sets.

Proof. Notice that, since the projections $S_{xy}(A) \to S_y(A)$ are closed[8], the following subsets of $S_y(A)$ are closed

$$A_n(y) = \text{“}\exists (a_i)_{i<n} \vDash p^{(n)} \restriction A \left(\vDash \bigwedge_{i<n-1} \varphi(a_i; y) \triangle \varphi(a_{i+1}; y) \right) \wedge \left(\vDash \varphi(a_{n-1}; y) \right)\text{”}$$

$$B_n(y) = \text{“}\exists (\tilde{a}_i)_{i<n} \vDash p^{(n)} \restriction A \left(\vDash \bigwedge_{i<n-1} \varphi(\tilde{a}_i; y) \triangle \varphi(\tilde{a}_{i+1}; y) \right) \wedge \left(\vDash \neg\varphi(\tilde{a}_{n-1}; y) \right)\text{”}$$

By Lemma 3.33 $d_p(x)\varphi(x; y) = \bigvee_{n \leq \mathrm{alt}(\varphi(x;y))} A_n(y) \wedge \neg B_{n+1}(y)$. □

Corollary 3.35 (T NIP)**.** Every $p \in S^{\mathrm{inv}}(\mathfrak{U}, M)$ is M-Borel-definable, i.e. every $d_{p(x)}\varphi(x; y)$ is a Borel subset of $S_y(M)$.

Proof. Constructible sets are Borel. □

Characterization of Forking

In NIP theories, forking can be characterized in terms of invariance.

Proposition 3.36 ([Sim15, Proposition 5.21], T NIP)**.** The converse of Proposition 2.78 holds, i.e. for global types A-non-forking is the same as Lstp_A-invariance.

Proof. Let $p \in S_x(\mathfrak{U})$ and $a_0 \equiv_{\mathrm{Lstp}_A} a_1$, and suppose that $\varphi(x; a_0) \triangle \varphi(x; a_1) \in p$. Up to replacing them with some $c_{j-1} \equiv_{\mathrm{Lstp}_A} c_j$ given by Corollary 2.43 we can assume[9] that they start an A-indiscernible $(a_i)_{i<\omega}$. Since $((a_{2i}, a_{2i+1}))_{i<\omega}$ is still A-indiscernible, $\pi(x) = \{\varphi(x; a_{2i}) \triangle \varphi(x; a_{2i+1}) \mid i < \omega\}$ is inconsistent: otherwise any $b \vDash \pi(x)$, together with $(a_i)_{i<\omega}$, would witness that the alternation rank of $\varphi^{\mathrm{opp}}(y; x)$ is infinite, contradicting NIP. But this means that the formula $\varphi(x; a_0) \triangle \varphi(x; a_1) \in p$ divides over A. □

Together with Remark 2.44, this has two immediate consequences:

Corollary 3.37 (T NIP)**.** A global type does not fork over A if and only if for all $M \supseteq A$ it is M-invariant.

Theorem 3.38 ([Sim15, Corollary 5.22], T NIP)**.** A global type does not fork over M if and only if it is M-invariant.

Remark 3.39 (T NIP)**.** By Corollary 3.37, for a global type, being non-forking over a small set is equivalent to being invariant over a small model.

[8] Inasmuch continuous from a compact to an Hausdorff.

[9] If $\varphi(x; c_{j-1}) \leftrightarrow \varphi(x; c_j) \in p$ always holds then $\varphi(x; a_0) \leftrightarrow \varphi(x; a_1) \in p$.

3.3 Measures

Dependency has quite strong repercussions on the behaviour of measures; this is basically due to the fact that, using the NIP, measures can be approximated by averages of types even if they are not smooth, and this allows to transfer results and constructions from $S(\mathfrak{U})$ to $\mathfrak{M}(\mathfrak{U})$. In the present section we prove this approximation result via the existence of smooth extensions and see its first "transfer" consequences.

Before we get into lemmas and theorems, let us see an example of how the finiteness of $\mathrm{alt}(\varphi(x;y))$ allows to build measures from indiscernible sequences.

Example 3.40 ([Sim15, Example 7.2])**.** Let $(a_i)_{i\in[0,1]}$ be any $\emptyset$-indiscernible sequence in a NIP theory indexed on the real unit interval, and let λ be the usual Lebesgue measure. Since for all $\varphi(x;b) \in L(\mathfrak{U})$ the NIP implies that $\mathrm{alt}(\varphi(x;b),(a_i)_{i\in[0,1]})$ is finite, $\{i \in [0,1] \mid\models \varphi(a_i;b)\}$ is a finite union of convex sets, hence Lebesgue-measurable. We can then define $\mu \in \mathfrak{M}_x(\mathfrak{U})$ setting $\mu(\varphi(x;b)) = \lambda(\{i \in [0,1] \mid\models \varphi(a_i;b)\})$.

Characterization of Forking

Proposition 3.41 ([Sim15, Proposition 7.15], T NIP)**.** Let A be small and $\mu \in \mathfrak{M}_x(\mathfrak{U})$. Then the following are equivalent.

1. is strongly Lstp A-invariant
2. μ is Lstp A-invariant
3. μ does not fork over A

Proof. One implication is obvious and another one is Corollary 2.117, so suppose that μ does not fork over A. Given $\varphi(x;y) \in L$ and $b_0 \equiv_{\mathrm{Lstp}_A} b_1$, we have to show that $\mu(\varphi(x;b_0) \triangle \varphi(x;b_1)) = 0$, and by hypothesis we only need to prove that $\varphi(x;b_0) \triangle \varphi(x;b_1))$ forks over A. If $((b_{2i}, b_{2i+1}) \mid i < \omega)$ is any indiscernible sequence starting with (b_0, b_1), then $\{\varphi(x;b_{2i}) \triangle \varphi(x;b_{2i+1}) \mid i < \omega\}$ is inconsistent by NIP, so this proves the proposition in the case where b_0, b_1 starts an A-indiscernible sequence. We can then conclude by Corollary 2.43 and the fact that $\mu(X \triangle Y)$ is a distance. □

Corollary 3.42 (T NIP)**.** Let M be small and $\mu \in \mathfrak{M}_x(\mathfrak{U})$. Then the following are equivalent.

1. is strongly M-invariant
2. μ is M-invariant
3. μ does not fork over M

Smooth Extensions

Lemma 3.43 ([Sim15, Lemma 7.6], T NIP)**.** Let $\mu \in \mathfrak{M}_x(M)$, $(b_i)_{i<\omega} \in M$ and $\varphi(x;y) \in L$. Then there is no $\varepsilon > 0$ such that for all $i < j < \omega$ we have $\mu(\varphi(x;b_i) \triangle \varphi(x;b_j)) > \varepsilon$.

Proof. First of all, applying the Standard Lemma inside $\widetilde{M}_\mu$, we can replace $(b_i)_{i<\omega}$ with some $L(\widetilde{M}_\mu)$-indiscernible $(\tilde{b}_i)_{i<\omega}$ in some elementary extension[10] of M. Hence, by NIP, $\pi(x) = \{\varphi(x;\tilde{b}_{2i}) \triangle \varphi(x;\tilde{b}_{2i+1}) \mid i < \omega\}$ must be inconsistent, or we would violate finite alternations. Then the existence of such an $\varepsilon > 0$ would trigger Proposition 2.116 and make $\pi(x)$ consistent. □

Hence, in a NIP theory, changing parameters infinitely many times inside a formula will result in two instances of it having small difference with respect to μ. This implies the existence of smooth measures with a prescribed restriction:

Proposition 3.44 ([Sim15, Proposition 7.9], T NIP)**.** Measures always have smooth extensions. More precisely, if $\mu \in \mathfrak{M}_x(M)$, then there is $N \succ M$ and an extension of M to $\mathfrak{M}_x(N)$ which is smooth.

Proof. The strategy of proof is to set up an iterative construction that must either stop and return the desired smooth extension or contradict the previous lemma. So suppose that μ has no smooth extension to any bigger model, and define $((M_\alpha, \mu_\alpha) \mid \alpha < |T|^+, \mu_\alpha \in \mathfrak{M}_x(M_\alpha))$ as follows. Start with $(M_0, \mu_0) = (M, \mu)$, if α is a non-zero limit take unions over the previous ordinals and at successor stages do the following. Since μ_α is not smooth by hypothesis, there are $\nu_0, \nu_1 \in L(\mathfrak{U})$ both extending μ_α and such that for some $\varphi_\alpha(x;y) \in L$, $b_\alpha \in \mathfrak{U}$ and a positive $\varepsilon_\alpha \in \mathbb{Q}$ we have $\nu_1(\varphi_\alpha(x;b_\alpha)) - \nu_0(\varphi_\alpha(x;b_\alpha)) > 4\varepsilon_\alpha$. Let $M_{\alpha+1}$ be a small model containing M_α and b_α and set $\mu_{\alpha+1} = \frac{\nu_0+\nu_1}{2} \restriction M_{\alpha+1}$.

Claim. For all $\theta(x) \in L(M_\alpha)$ we have $\mu_{\alpha+1}(\theta(x) \triangle \varphi_\alpha(x;b_\alpha)) > \varepsilon_\alpha$.

Proof of the Claim. Denote $\varphi_\alpha(x;b_\alpha)$ with φ and $\theta(x)$ with θ. If we show that at least one between $\nu_0(\varphi \triangle \theta)$ and $\nu_1(\varphi \triangle \theta)$ is greater than $2\varepsilon_\alpha$ we are done, so assume that this is not the case. Since $\theta \in L(M_\alpha)$ and both ν_0 and ν_1 extend μ_α, we have $\nu_1(\theta) = \nu_0(\theta)$. But then[11] we can get an absurd as follows:

$$\begin{aligned}
&\nu_1(\varphi \triangle \theta) + \nu_0(\varphi \triangle \theta) \leq 4\varepsilon_\alpha < \nu_1(\varphi) - \nu_0(\varphi)\\
\Longrightarrow&\nu_1(\varphi) + \nu_1(\theta) - 2\nu_1(\varphi \wedge \theta) + \nu_0(\varphi) + \nu_0(\theta) - 2\nu_0(\varphi \wedge \theta) < \nu_1(\varphi) - \nu_0(\varphi)\\
\Longrightarrow&2\big(\nu_1(\theta) + \nu_0(\varphi)\big) < 2\big(\nu_1(\varphi \wedge \theta) + \nu_0(\varphi \wedge \theta)\big)\\
\Longrightarrow&\nu_1(\theta) + \nu_0(\varphi) < \nu_1(\varphi \wedge \theta) + \nu_0(\varphi \wedge \theta) \leq \nu_1(\theta) + \nu_0(\varphi) \Longrightarrow 0 < 0
\end{aligned}$$

□ CLAIM

[10] After these $\tilde{b}_i$ are found in some $\widetilde{N} \succ \widetilde{M}_\mu$, just take the reduct $\widetilde{N} \restriction L$ of $\widetilde{N}$ to L.

[11] The fact that $\nu_1(\theta) = \nu_0(\theta)$ is used when passing from the second line to the third one.

Carry out this construction for all $\alpha < |T|^+$. By pigeonhole, up to discarding co-$|T|^+$-many α, we can assume that $\varphi_\alpha(x;y)$ and ε_α do not depend on α and remove the subscript. But if $\beta < \alpha$ then $\varphi(x;b_\beta)$ counts as a θ in the Claim, hence $\bigcup_{\alpha<|T|^+}\mu_\alpha$ gives measure at least ε to all $\varphi(x;b_\beta) \triangle \varphi(x;b_\alpha)$, against Lemma 3.43. □

Approximation of Measures

When dealing with types, one can consider its realizations. Strictly speaking (and without resorting to continuous logic) measures which are not types cannot be realized but, for instance, if μ is an average of types we still have their realizations available. Proposition 2.115 tells us that, in any theory, if μ is smooth we can get close enough to this scenario. In NIP theories, we can approximate with averages of types even non-smooth measures.

Theorem 3.45 ([Sim15, Proposition 7.11], T NIP)**.** Let A be small and $\mu \in \mathfrak{M}_x(A)$. For all $B_0, \dots, B_{m-1}$ Borel subsets of $S_x(A)$, $\varphi(x;y) \in L$ and $\varepsilon > 0$ there are $p_0, \dots, p_{n-1} \in S_x(A)$ such that for all $b \in A$ and $j < m$

$$|\mu(B_j \cap \varphi(x;b)) - \operatorname{Av}(p_0, \dots, p_{n-1}; B_j \cap \varphi(x;b))| \le \varepsilon$$

Moreover, we can choose such $p_0, \dots, p_{n-1}$ in $S(\mu)$.

Proof. By Proposition 3.44, let $M \supseteq A$ be such that there is a smooth $\nu \in \mathfrak{M}_x(M)$ extending μ. It then suffices to apply Proposition 2.115 to the unique extension $\nu \mid \mathfrak{U}$ of ν to $\mathfrak{U}$, set $p_i = \operatorname{tp}(a_i/A)$, and notice that since all the $\operatorname{tp}(a_i/\mathfrak{U})$ are in the support of $\nu \mid \mathfrak{U}$, as a special case if $\varphi(x) \in \operatorname{tp}(a_i/A)$ then $\mu(\varphi(x)) = \nu(\varphi(x)) > 0$, hence every p_i is in $S(\mu)$. □

We can now take advantage of approximations to transfer results from types to measures.

Theorem 3.46 ([Sim15, Proposition 7.19], T NIP)**.** Let M be small. Then M-invariant global measures are M-Borel-definable.

Proof. Given $\mu \in \mathfrak{M}^{\mathrm{inv}}(\mathfrak{U}, M)$, $\varphi(x;y) \in L$ and C closed in $[0,1]$ we have to show that $X = \{q \mid \exists/\forall b \in q(\mathfrak{U})\ \mu(\varphi(x;b)) \in C\}$ is Borel in $S_y(M)$. For all $s \in \mathbb{Q} \cap [0,1] \cap C^{\complement}$ choose ε_s such that $(s-\varepsilon_s, s+\varepsilon_s) \cap [0,1] \subseteq [0,1] \setminus C$, and notice that we may suppose that for all k cofinitely many ε_s are smaller than 2^{-k}. Apply[12] Theorem 3.45 to μ, $\varphi(x;y)$ and ε_s, choosing $p_0, \dots, p_{n-1}$ in the support of μ such that for all $b \in \mathfrak{U}$

$$|\mu(\varphi(x;b)) - \operatorname{Av}(p_0, \dots, p_{n-1}; \varphi(x;b))| \le \varepsilon_s$$

[12] The cautious reader may have noticed that $\mathfrak{U}$ is not small, so we have to work in a bigger $\widetilde{\mathfrak{U}} \,{}^+\!\succ \mathfrak{U}$.

and the set $\nu_s(x) = \mathrm{Av}(p_0, \dots, p_{n-1})$. Since by Corollary 3.42 μ is *strongly* M-invariant, by Proposition 2.104 the p_i are M-invariant, so by Corollary 3.35 $Y_s = \{q \in S_y(M) \mid \exists/\forall b \in q(\mathfrak{U})\ \nu_s(\varphi(x;b)) = s\}$ is Borel. Therefore it suffices to show that

$$X = \bigcap\{Y_s^{\complement} \mid s \in \mathbb{Q} \cap [0,1] \cap C^{\complement}\}$$

Let $b \vDash q$. If there is s such that $q \in Y_s$, then by construction $\mu(\varphi(x;b)) \in [0,1] \setminus C$, so $q \notin X$. Conversely, if $\mu(\varphi(x;b)) = r \in [r-\varepsilon, r+\varepsilon] \cap [0,1] \subseteq [0,1] \setminus C$, we can find an $s \in [r-\varepsilon, r+\varepsilon] \cap [0,1] \cap \mathbb{Q}$ such that $\varepsilon_s < \varepsilon$ because this does not happen only for finitely many s. Then we have $\nu_s(\varphi(x;b)) = s \in [r-\varepsilon_s, r+\varepsilon_s]$ and therefore $q \in Y_s$. □

Remark 3.47. Notice that in the previous proof we used a countable intersection, so even if the Y_s are constructible we have proved that X is Borel, but not that X is constructible.

Although we will use Theorem 3.45 again later in this thesis, we will be forced by space reasons to omit a large number of its consequences. For instance, a beautiful fact proved in [CPS14, Theorem 2.7] is that in NIP theories definability of types implies definability of measures, and the proof relies heavily on being able to approximate measures with averages of types.

Borel-definability of invariant measures implies that it is possible to define a *product measure* $\mu \otimes \nu$, and a lot of results generalize from the product of invariant types to the product of invariant measures. Notice that the usual measure-theoretic product will not suffice, since it would only account for formulas in which the variables in x are separated from the ones in y. In other words, the space $S_{xy}(\mathfrak{U})$ is not the product of $S_x(\mathfrak{U})$ and $S_y(\mathfrak{U})$: in the latter there is no place for formulas like $x = y$. This is analogous to the fact that the Zariski topology on $\mathbb{A}^{n+m}$ is not the product topology on $\mathbb{A}^n \times \mathbb{A}^m$. Since we will not need product measures, we are not going to see these results, and refer the reader to [Sim15, Chapter 7].

3.4 Honest Definitions

In stable theories, externally definable sets of any model are already definable and sets are *stably embedded*. This means that, for any A and $\varphi \in L(\mathfrak{U})$, there is some $\psi \in L(A)$ such that $\varphi(A) = \psi(A)$. Indeed, this is another characterization of stability (see [TZ12, Exercise 10.1.5]), so in unstable NIP theories it may fail. We will now study a result from [CS13] implying that, in NIP theories, the parameters of ψ may be chosen in an "elementary extension" of A. It will also imply that such a ψ may be required to have a certain property, called *honesty*, that will be used to show quantifier elimination for M^{ext}.

Definition and Existence

Definition 3.48. If M is an L-structure and $A \subseteq M$, the *pair* (M, A) is defined as follows. Let P be a unary predicate outside L, and define $L_P = L \cup \{P\}$. Then let (M, A) be the L_P-expansion of M obtained by interpreting P with A. An *elementary pair* is a pair (N, M) with $M \prec_L N$.

Remark 3.49. Even if M and N are NIP and $M \prec N$, the pair (N, M) may be independent. As Poizat points out in [Poi83, Section 3], this can be the case even with pairs of stable structures.

Definition 3.50. Consider a pair (M, A), a formula $\varphi(x; y) \in L$ and $b \in M$. An *honest definition of* $\varphi(x; b)$ *over* A is given by an elementary extension $(M', A') \succ (M, A)$ and a formula $\psi(x; d)$ such that

- $\varphi(x; z) \in L$ and $d \in A'$
- $\varphi(A; b) = \psi(A; d)$
- $\psi(A'; d) \subseteq \varphi(A'; b)$ (*honesty condition* or *honesty hypothesis*)

Example 3.51 (Dishonest definition in DLO). Let $A = \mathbb{Q}$, $M \supseteq \mathbb{Q} \cup \{\pi\}$ and $(M', A') \ ^{+}\!\!\succ (M, A)$. By $\aleph_1$-saturation find $\pi + \varepsilon \in A'$ such that $\pi + \varepsilon > \pi$ and for all $m \in \mathbb{Q}$ such that $m > \pi$ we have $\pi + \varepsilon < m$. Then $x \leq \pi + \varepsilon$ is *not* an honest definition of $x < \pi$, even if the two formulas agree on $A = \mathbb{Q}$.

We state the following result in a way that is even too general for the purposes of the present section, and indeed the exposition in [Sim15, Section 3.2] works well with a special case of it. Anyway, the greater generality will be needed later on. We first need a definition.

Definition 3.52. Let $\varphi(x; y)$ be a formula and $p(x)$ a partial type. We say that φ *is* NIP *over* $p(x)$ iff for all $c \in \mathfrak{U}$ there is no indiscernible sequence $(a_i)_{i<\omega}$ *of realizations of* p such that $\operatorname{alt}(\varphi(x; c), (a_i)_{i \in I}) = \infty$.

Lemma 3.53 ([CS13, Proposition 1.1]). Let $L \subseteq L'$ and work in a monster $\mathfrak{U}$ that is an L'-structure. Let $p(x)$ be a small partial L'-type, $\varphi(x; y) \in L$ a NIP formula over $p(x)$ and $c \in \mathfrak{U}$. Then for each small $A \subseteq p(\mathfrak{U})$ there is $\theta(x) \in L(p(\mathfrak{U}))$ such that

1. $\theta(x) \cap A = \varphi(x; c) \cap A$
2. $\theta(p(\mathfrak{U})) \subseteq \varphi(p(\mathfrak{U}); c)$
3. $\varphi(x; c) \setminus \theta(x)$ does not contain any A-invariant global L-type consistent with $p(x)$.

Proof. Let X be the compact[13] space of all $q(x) \in S^{\mathrm{inv}}_x(\mathfrak{U} \upharpoonright L, A)$ consistent with $\{\varphi(x;c)\} \cup p(x)$. Suppose that for all $q \in X$ we are able to find some $\psi \in L(p(\mathfrak{U}))$ such that $p(x) \vdash \psi(x) \to \varphi(x;c)$. Then we can write $X = \bigcup_{q\in X}[\psi_q]$, extract a finite subcover by compactness, and let θ be its disjunction. Such a θ will trivially satisfy the last two points of the thesis, and the fact that it satisfies the first one follows because among the A-invariant global L-types there are the L-types of points in A.

So fix $q \in X$. We set up an iterative construction that will stop by NIP and produce the required ψ. The construction finds inductively $a_i, b_i \in p(\mathfrak{U})$ and $q_i \subseteq q_{i+1} \subseteq q$ in the following way

At stage $4i + 0$ *set* $q_{2i}(x) = q(x) \upharpoonright Aa_{<i}, b_{<i}$

At stage $4i + 1$ *let* $a_i \models q_{2i}(x) \cup \{\varphi(x;c)\} \cup p(x)$

At stage $4i + 2$ *set* $q_{2i+1}(x) = q(x) \upharpoonright Aa_{\le i}, b_{<i}$

At stage $4i + 3$ *let* $b_i \models q_{2i+1}(x) \cup \{\neg\varphi(x;c)\} \cup p(x)$

By Corollary 2.26, if we are able to carry out this construction for all the $i < \omega$, then we will find inside $p(\mathfrak{U})$ a Morley sequence of q over A, which in particular will be A-indiscernible, alternating truth values on $\varphi(x;c)$ infinitely many times and contradicting the fact that $\varphi(x;y)$ is NIP over $p(x)$. So, for some i, the construction must stop. Stages with even index cannot be stopped, since nothing can prevent us from taking a restriction. Neither can $4i+1$, since the small type $q_{2i} \cup \{\varphi(x;c)\} \cup p$ is included in $q \cup \{\varphi(x;c)\} \cup p$, which is guaranteed to be consistent by hypothesis. Therefore the construction stops because there is no b_i to fulfil the requirements of stage $4i+3$. This means that $p \cup q_{2i+1} \vdash \varphi(x;c)$, and by compactness we can find the required ψ inside $q_{2i+1} \subseteq q$. Since q_{2i+1} has parameters from $Aa_{\le i}, b_{<i} \subseteq p(\mathfrak{U})$, we are done. □

Theorem 3.54 ([CS13, Corollary 1.3], [Sim15, Theorem 3.13], T NIP)**.** Let $A \subseteq M$, $\varphi(x;y) \in L$ and $b \in M$. Then any $(M', A') \mathrel{^{+}\!\succ} (M, A)$ has an honest definition of $\varphi(x;b)$ over A.

Proof. Apply Lemma 3.53 setting

$$L' := L_P \quad \mathfrak{U} := (M', A') \quad p := \{P\} \quad \varphi := \varphi \quad c := b \quad A := A$$

Since $\theta \in L(p(\mathfrak{U})) = L(P(M', A')) = L(A')$, we can let θ be the required honest definition. □

For a version of the previous theorem that does not mention any (M', A') see [Sim15, Remark 3.14].

[13] It is a closed subspace of the space of L'-types $S_x(\mathfrak{U})$.

Quantifier Elimination for Shelah's Expansion

Honest definitions "stay in the M-closure":

Proposition 3.55 ([Sim15, Proposition 3.21], T NIP)**.** All externally definable $D \subseteq M^k$ have an external definition $\varphi(x; b)$ with the following property: if $\theta(x; a) \in L(M)$ is such that $D \subseteq \theta(M; a)$, then $\mathfrak{U} \vDash \varphi(x; b) \to \theta(x; a)$.

Proof. Let $N \ {}^{+}\!\succ M$, consider the pair (N, M) and let $\varphi_0(x; b_0) \in L(N)$ be such that $\varphi_0(M; b_0) = D$. Let $\varphi(x; b)$ be an honest definition of $\varphi_0(x; b_0)$ in some $(N', M') \ {}^{+}\!\succ (N, M)$. If $D \subseteq \theta(M; a)$ then we have $(N, M) \vDash \forall x \in P \ \varphi_0(x; b_0) \to \theta(x; a)$. Since by honesty $\varphi(M'; b) \subseteq \varphi_0(M'; b_0)$, then $\varphi(M'; b) \subseteq \theta(M'; a)$. Since M' is a model this means $\vDash \forall x \ \varphi(x; b) \to \theta(x; a)$. □

Theorem 3.56 ([Sim15, Proposition 3.23], T NIP)**.** If M is NIP, then $\mathrm{Th}(M^{\mathrm{ext}})$ has quantifier elimination.

Proof. By induction on formulas, all we have to show is that if $D \subseteq M^{k_0+k_1}$ is externally definable and $\pi \colon M^{k_0+k_1} \to M^{k_0}$ is the usual projection, then $\pi(D)$ is still externally definable.

Let $\varphi(x_0, x_1; b)$ be an external definition of D as given by Proposition 3.55 with $b \in M' \ {}^{+}\!\succ M$, use M' to build L^{ext} and set $\psi(x_0; b) = \exists x_1 \ \varphi(x_0, x_1; b)$. Since $\pi(D) \equiv \exists x_1 \in M \ \varphi(x_0, x_1; b)$, we have $\pi(D) \subseteq \psi(M; b)$. If we prove the other inclusion, we can conclude that

$$M^{\mathrm{ext}} \vDash \exists x_1 R_{\varphi(x_0,x_1;b)}(x_0, x_1) \leftrightarrow R_{\psi(x_0;b)}(x_0)$$

Let $a \in M^{k_0} \setminus \pi(D)$ and set $\zeta(x_0, x_1; a) = x_0 \neq a$. Since $\pi(D) \subseteq \zeta(M)$, by choice of φ we have $\vDash \varphi(x_0, x_1; b) \to \zeta(x_0, x_1; a)$. But then $\vDash \psi(x_0; b) \to x_0 \neq a$, and so $\psi(M; b) \subseteq \pi(D)$. □

Corollary 3.57 ([Sim15, Corollary 3.24])**.** If M is NIP, then M^{ext} is NIP too.

Proof. By quantifier elimination and the fact that R commutes with connectives[14] it suffices to check formulas of the form $R_{\varphi(x;y)}(x; y)$. Since every $L(M^{\mathrm{ext}})$-indiscernible sequence is $L(M)$-indiscernible, $R_{\varphi(x;y)}(x; y)$ has at most the same alternation rank as $\varphi(x; y)$. □

Corollary 3.58. If M is NIP, then $M^{\mathrm{ext}^{\mathrm{ext}}}$ can be identified with M^{ext}. In other words, all externally definable subsets of M^{ext} are definable.

Proof. Let $N_0 \ {}^{+}\!\succ M$ and use N_0 to build the language L^{ext}; then let $\widetilde{N} \ {}^{+}\!\succ M^{\mathrm{ext}}$. By Theorem 3.56, every L^{ext} formula is equivalent to some

[14] Or by Lemma 3.24.

$R_{\varphi(x;y,b)}(x;y)$ for some $\varphi(x;y,z) \in L$ and $b \in N_0$. Let $a \in \widetilde{N}$ and consider $R_{\varphi(x;y,b)}(x;a)$. By saturation, we can find inside N_0 some (a_0, b_0) with the same L-type over M as (a,b). Then for all $m \in M$ we have $\models \varphi(m;a,b) \leftrightarrow \varphi(m;a_0,b_0)$. Since M and M^{ext} share the same underlying set, writing $m \in M$ or $m \in M^{\text{ext}}$ does not make any difference. Therefore for all $m \in M^{\text{ext}}$ we have $\widetilde{N} \models R_{\varphi(x;y,b)}(m;a) \leftrightarrow R_{\varphi(x;a_0,b_0)}(m)$, and the externally definable set $R_{\varphi(x;y,b)}(M^{\text{ext}};a)$ can be defined inside M^{ext} by the L^{ext}-formula $R_{\varphi(x;a_0,b_0)}(x)$. □

3.5 The Type-Connected Component

This section revolves around two main results concerning the behaviour of type-definable, bounded-index subgroups of a definable group in NIP theories. The first one, Theorem 3.64, ensures that G^{00}_A does not depend on A, while in Corollary 3.78 we will prove that passing from M to M^{ext} does not change G^{00}.

Existence

We will now prove that in NIP theories G^{00} exists[15]. Nonetheless, our first lemmas are still valid in any theory.

Lemma 3.59. If A is small and X is a type-definable set invariant under $\text{Aut}(\mathfrak{U}/A)$, then it is A-type-definable.

Proof. We want to write X as the intersection of a family of $L(A)$-definable sets. The natural candidate for such a family is clearly $\Phi = \{\varphi(x) \in L(A) \mid X \subseteq [\varphi(x)]\}$. Of course $X \subseteq \bigcap \Phi$; suppose that the other inclusion does not hold, and let b witness it. If we find $c \in X$ such that $c \equiv_A b$ then we are done: sending c to b by an element of $\text{Aut}(\mathfrak{U}/A)$ contradicts the fact that X should be fixed set-wise by such an automorphism. Suppose there is no such c; then by compactness, since X is type-definable, there is $\psi(x) \in \text{tp}(b/A)$ such that $X \subseteq [\neg\psi(x)]$. Then $\neg\psi \in \Phi$, so $\models \neg\psi(b)$, a contradiction. □

Notation 3.60. *G will indicate an $\emptyset$-definable group.*

Lemma 3.61 ([BOPP05, Remark 1.4])**.** Every type-definable subgroup $H < G$ can be written as $\bigcap_{i \in I} H_i$, where each H_i is a subgroup of G that can be type-defined by countably many formulas.

Proof. Let $H = \bigcap_{j \in J}[\varphi_j(x)]$, where every φ_j concentrates on G. Suppose without loss of generality that $\models \varphi_j(x) \to \varphi_j(x^{-1})$ and fix $j_0 \in J$. Iteratively, for all $k \in \omega$ find by compactness $j_{k+1} \in J$ such that $\models (\varphi_{j_{k+1}}(x) \wedge \varphi_{j_{k+1}}(y)) \to \varphi_{j_k}(x \cdot y)$, and let $H_{i_0} = \bigcap_{k \in \omega}[\varphi_{j_k}]$, which is clearly a subgroup

[15] See Definition 1.81.

of G. Now let $j_\omega \in J \setminus \{j_k \mid k < \omega\}$ and repeat the construction to find H_{i_1}. After less than $|J|^+$ steps the process ends and produces the required subgroups H_i. □

This is the crucial result:

Proposition 3.62 ([HPP08, Proposition 6.1]). Suppose that $\bar{a} = (a^\ell \mid \ell < \omega) \in \mathfrak{U}$ and $H = \Sigma(x; \bar{a})$, where $\Sigma(x; \bar{y}) = \{\varphi_n(x; \bar{y}) \mid n < \omega\} \subseteq L$. If H is a bounded-index subgroup of G, and each $\varphi_n(x; \bar{y})$ is NIP, then its orbit under the action of $\mathrm{Aut}(\mathfrak{U}/\emptyset)$ is bounded.

Proof. Observe that, since Σ is countable, up to conjunctions we can assume $\varphi_n \to \varphi_{n-1}$. Suppose that H has unbounded orbit. The proof proceeds in four steps, corresponding to paragraphs.

By hypothesis we can find unboundedly many distinct $\Sigma(x; \bar{b}_i)$. Then, by Theorem A.49, we can find an indiscernible sequence $(\bar{a}_i \mid i < \omega)$ such that the $H_i = \Sigma(x; \bar{a}_i)$ are pairwise distinct.

Suppose that for some $i_0 < \omega$ we have $H_{i_0} \supseteq \bigcap_{j \neq i_0} H_j$. Then, using the Standard Lemma, we can replace $\bar{a}_{i_0}$ with an indiscernible sequence of unbounded length between $\bar{a}_{i_0-1}$ and $\bar{a}_{i_0+1}$ in a way that the resulting long sequence is still A-indiscernible and with the same Ehrenfeucht-Mostowski type. Since the H_i were pairwise distinct, this would produce too many subgroups $K \supseteq \bigcap_{j \neq i_0} H_j$, contradicting the fact that the latter has bounded index, since it is a bounded intersection of bounded index subgroups. We can therefore choose, for each $i < \omega$, some $c_i \in \bigcap_{j \neq i} H_j \setminus H_i$.

We now use the Standard Lemma and automorphisms to extend $(\bar{a}_i \mid i < \omega)$ to an indiscernible sequence of sufficiently big length and choose $c_i \in H_i$ as before. This allows us to apply Theorem A.49 again and assume that $((c_i, \bar{a}_i) \mid i < \omega)$ is indiscernible.

Now we want to show that some $\varphi_n(x; y)$ has the IP, contradicting the hypotheses. We will do this by finding some $n \in \omega$ and, for each $w \in \mathscr{P}_{\mathrm{fin}}(\omega)$, some d_w such that $\neg\varphi_n(d_w; \bar{a}_i) \iff i \in w$. If we take $d_w = \prod_{i \in w} c_i$, then if $i \notin w$ we have $d_w \in H_i$, so $\vDash \varphi_n(d_w; \bar{a}_i)$ for all n. Thus we are left to find an n for which the other implication holds. In order for such an n to work, it is sufficient that for all $x_0, x_1 \in H_i$ we have $\vDash \neg\varphi_n(x_0 c_i x_1; \bar{a}_i)$. Since $x_0 c_i x_1 \notin H_i$ and each φ_k implies the previous ones[16], by compactness there is some n_i such that $x_0 \in H_i \wedge x_1 \in H_i \vdash \neg\varphi_{n_i}(x_0 c_i x_1; \bar{a}_i)$. Since $((c_i, \bar{a}_i) \mid i < \omega)$ is indiscernible, n_i does not depend on i. □

Remark 3.63. The NIP assumption was only used in the last step, when we got an absurd in showing that one of the φ_n has the independence property. Therefore, as long as they are NIP, the result is valid even if T is independent. This will be useful later on in this section.

[16] Otherwise we would only have had $\bigvee_{i<\ell} \neg\varphi_{n_i}$.

Theorem 3.64. In NIP theories, G^{00} exists and has index bounded by $2^{|T|}$.

Proof. Given a small A, we have to show that, $G^{00}_A = G^{00}_\emptyset$. Clearly $\subseteq$ always holds. Write $G^{00}_A = \bigcap_{i \in I_A} H_i$ as in Lemma 3.61. Since each H_i has bounded index, inasmuch it has G^{00}_A as a subgroup, by Proposition 3.62, it has bounded orbit under $\mathrm{Aut}(\mathfrak{U}/\emptyset)$. Let λ be a common bound for $|\mathrm{Orb}(H_i)|$ as $i \in I_A$ varies, and write, for each H_i, $\mathrm{Orb}(H_i) = \{f_{i,j}(H_i) \mid j < \lambda\}$. Now consider

$$K_A = \bigcap_{\substack{j<\lambda \\ i \in I_A}} f_{i,j}(H_i)$$

Since K_A is trivially invariant under $\mathrm{Aut}(\mathfrak{U}/\emptyset)$, by Lemma 3.59 it is $\emptyset$-type-definable, and it has bounded index because it is a bounded[17] intersection of subgroups of bounded index. This means that we have[18]

$$G^{00}_\emptyset \supseteq G^{00}_A \supseteq K_A \supseteq G^{00}_\emptyset$$

and this proves that for all small A we have $G^{00}_\emptyset = G^{00}_A$. The bound on the index follows from Proposition 2.54 and Löwenheim-Skolem. □

Another name for G^{00} is *the infinitesimal subgroup*. For an explanation, see [HPP08, right before Proposition 6.2]. It was proven in [BOPP05] that if G is $\emptyset$-definable in an o-minimal theory, then G/G^{00} is a real Lie group. This was part of *Pillay's conjectures*, now theorems, relating o-minimal structures and Lie groups. See also [HPP08, Theorem 8.1].

Passing to Shelah's Expansion

If M is NIP, by Theorem 3.56 we can identify $S^{\mathrm{ext}}_G(M)$ with $S_G(M^{\mathrm{ext}})$ and carry out the Ellis group construction directly inside $S_G(M^{\mathrm{ext}})$. Nevertheless, dealing with externally definable sets still introduces some degree of complication. Hence we would be happy to assume from the start that all externally definable subsets are already definable. This can be done by replacing M with M^{ext}, and by Corollary 3.58 all the results from Theorem 1.75 to Theorem 1.79 hold setting $M := M^{\mathrm{ext}} =: N$. Moreover, the NIP hypothesis transfers from M to M^{ext} by Corollary 3.57. This settles the problem, at least for what concerns one of the two objects we want to compare, namely the Ellis group. Alas, since in L^{ext} there are more predicate symbols, hence potentially more bounded index $\emptyset$-type-definable subgroups to intersect, we have two questions to answer: if $\widetilde{\mathfrak{U}}$ is a monster for M^{ext} and $\mathfrak{U} = \widetilde{\mathfrak{U}} \restriction L$ is the corresponding monster for M, is $G^{00}(\widetilde{\mathfrak{U}}) = G^{00}(\mathfrak{U})$? If this is the case, is the logic topology on G/G^{00} unchanged? In [CPS14, Sections 4.2 and 4.3] it was shown that the answer to both questions is *yes*, and we are now going to explain why.

[17] $|I_A|$ is bounded by $(|A| + \aleph_0)^{\aleph_0}$.

[18] $G^{00}_A \supseteq K_A$ is simply due to the fact that we are taking a bigger intersection.

Notation 3.65. Until the end of Section 3.5 we assume $G(\mathfrak{U}) = \mathfrak{U}$ in order to simplify the notation. Moreover, "saturated" and "κ-saturated" will implicitly mean also "strongly homogeneous" and "κ-strongly homogeneous".

Hereditary Subgroups

The main proof will involve a certain inductive construction that must stop by NIP. This implies that we will need to prove some lemmas to set up said construction, and these will be quite technical. To help the reader see where this is going, we anticipate that the intended use of the following will be when $N \mathrel{{}^{+}{\succ}} M$ is used to define L^{ext}.

Definition 3.66. Work inside an elementary pair (N, M). Let $\Sigma(x)$ be a *disjunction* of complete L-types over some $A \subseteq N$, each consistent with $P(x)$. We call $\Sigma(x)$ an *hereditary subgroup* of $P(x)$ iff for all $(N', M') \succ (N, M)$ we have that $\Sigma(M')$ is a subgroup of M'. An hereditary subgroup is *of hereditarily bounded index* iff every time $(N', M') \succ (N, M)$ is κ-saturated and not κ^+-saturated then $[M' : \Sigma(M')] < \kappa$.

At least in the cases that will be relevant to us, these two notions have syntactic characterizations that allow to avoid checking all the $(N', M') \succ (N, M)$. This is the content of the next lemmas.

Lemma 3.67 ([CPS14, Lemma 4.7])**.** Work inside an elementary pair (N, M). Let $\Sigma(x)$ be a disjunction of complete L-types over some $A \subseteq N$, each consistent with $P(x)$. Consider the following two statements:

1. The following two conditions hold:
   (a) If $p(x) \in \Sigma(x)$ then $p(x^{-1}) \in \Sigma(x)$.
   (b) For all $p(x), q(x) \in \Sigma(x)$ and all sequences $\Phi = (\varphi_r(x) \mid r \in \Sigma)$ of formulas $\varphi_r(x) \in r(x)$ there are $\psi_p \in p$, $\psi_q \in q$, $n \in \omega$ and, for $i < n$, some $r_i \in \Sigma$, such that
   $$\big(\psi_p(x) \wedge P(x) \wedge \psi_q(y) \wedge P(y)\big) \to \bigvee_{i<n} \varphi_{r_i}(x \cdot y) \wedge P(x \cdot y)$$
2. $\Sigma(M)$ is a subgroup of M.

Then $1 \Rightarrow 2$ always holds. If moreover (N, M) is $|A|^+$-saturated, then $2 \Rightarrow 1$ holds too.

Proof. ($1 \Rightarrow 2$) Let $a, b \vDash \Sigma(x) \wedge P(x)$ and call $p(x) = \mathrm{tp}(a/A)$, $q(x) = \mathrm{tp}(b/A)$. Since $P(N, M) = M$ is a group and $\vDash P(a) \wedge P(b)$, then by (1a) we have $a^{-1}, b^{-1} \vDash \Sigma \wedge P$. Similarly, $\vDash P(a \cdot b)$. If $a \cdot b \notin \Sigma$, then for all $r \in \Sigma$ there is some $\varphi_r \in r$ such that $\vDash \neg\varphi_r(a \cdot b)$. This contradicts (1b).

(2 ⇒ 1) If $\Sigma(M)$ is a group, for any $m, n \in \Sigma(M)$ we have that for $p = \mathrm{tp}(m/A)$ and $q = \mathrm{tp}(n/A)$ both (1a) and (1b) hold. Therefore, the only possible obstruction is that some $p \in \Sigma$ is not realized in M. Since p is consistent with P, this is the same as saying that it is not realized in (N, M), and by $|A|^+$-saturation of the latter this cannot happen. □

Lemma 3.68 ([CPS14, Lemma 4.8])**.** Let (N, M) be κ-saturated and $A \subseteq N$ be such that $|A| \lll \kappa$. Let $\Sigma(x)$ be a disjunction of complete L-types over A, each consistent with $P(x)$. Suppose that $\Sigma(M)$ is a subgroup of M. Then the following are equivalent:

1. For every sequence $\Phi = (\varphi_r \mid r \in \Sigma)$ of formulas $\varphi_r \in r$, there are $n, m \in \omega$ and, for $k < n$, some $\varphi_k \in \Phi$, such that if $(a_i)_{i<m} \in M$ are pairwise distinct, then there are $i < j < m$ such that $a_i^{-1}a_j \vDash \bigvee_{k<n} \varphi_k(x)$.
2. $[M : \Sigma(M)] < \kappa$.

Proof. (1 ⇒ 2) Fix $\lambda \geq \kappa$. Since $[M : \Sigma(M)]$ is unbounded, we can find $\bar{a} = (a_i)_{i<\lambda}$ in M such that for all $i < j < \lambda$ we have $a_i^{-1}a_j \notin \Sigma(M)$. Hence for all $r \in \Sigma$ there is $\varphi_r^{i,j} \in r$ such that $a_i^{-1}a_j \vDash \neg\varphi_r^{i,j}$. Colour $\{i, j\} \in [\lambda]^2$ with $\Phi_{i,j} \colon r \mapsto \varphi_r^{i,j}$. Taking λ big enough, by Erdös-Rado[19], there is an infinite $I \subseteq \lambda$ where these maps are constantly $\Phi = (\varphi_r \mid r \in \Sigma)$. Then, for all $i, j \in I$, we have $a_i^{-1}a_j \vDash \bigwedge_{r\in\Sigma} \neg\varphi_r(x)$, contradicting 1.

(2 ⇒ 1) Suppose that $[M : \Sigma(M)] = \lambda \ll \kappa$, but 1 does not hold for Φ. By compactness and saturation, we can find $(a_i \mid i < \lambda^+)$ inside N such that $a_i^{-1}a_j \vDash \bigwedge_{r\in\Sigma} \neg\varphi_r(x)$, and this means that $[M : \Sigma(M)] \geq \lambda^+$. □

Notice how the previous lemma allows to speak of "bounded index" without mentioning any bound, analogously to Proposition 2.39.

Corollary 3.69 ([CPS14, Remark 4.9])**.** If (N, M) is sufficiently saturated[20], $\Sigma(x)$ is an hereditary subgroup of $P(x)$ of hereditarily bounded index if and only if satisfies point 1 of both the previous lemmas.

Proof. After possibly replacing M with $P(N, M)$ in the statements of such points, they hold in (N, M) if and only if they hold in all $(N', M') \succ (N, M)$. □

Since we are going to do an inductive construction, we will need these notions to be preserved under unions of elementary chains. This will require some hypotheses.

Notation 3.70. $L_P^{\forall,\mathrm{bdd}} = \{\forall z \in P\ \varphi(z; \bar{y}) \mid \varphi \in L\}$

[19]Theorem A.1

[20]This condition can be removed if one also removes "of hereditarily bounded index" and replaces "both the previous lemmas" with "Lemma 3.67".

Lemma 3.71 ([CPS14, Lemma 4.10])**.** Let $\Sigma(x;\bar{y})$ a disjunction of complete L-types, each consistent with $P(x)$. Then the following facts hold:

1. If (N_0, M_0), (N_1, M_1) are saturated elementary pairs, $\bar{b}_i \in N_i$ and $\mathrm{tp}(\bar{b}_0/L_P^{\forall,\mathrm{bdd}}) = \mathrm{tp}(\bar{b}_1/L_P^{\forall,\mathrm{bdd}})$, then $\Sigma(x;\bar{b}_0)$ is an hereditary subgroup of $P(x)$ of hereditarily bounded index in (N_0, M_0) if and only if the same is true for $\Sigma(x;\bar{b}_1)$ in (N_1, M_1).
2. Suppose that $(N_i \mid i < \kappa)$ and $(M_i \mid i < \kappa)$ are elementary chains of L-structures, that (N_i, M_i) is an elementary pair for all $i < \kappa$, and let $N = \bigcup_{i<\kappa} N_i$ and $M = \bigcup_{i<\kappa} M_i$. If $\bar{b} \in N_0$ is such that $\mathrm{tp}(\bar{b}/L_P^{\forall,\mathrm{bdd}})$ is the same in all the (N_i, M_i), then it is still unchanged in (N, M).
3. As a special case, if $\Sigma(x;\bar{b})$ is an hereditary subgroup of $P(x)$ of hereditarily bounded index in (N_0, M_0), then the same is true in (N, M).

Proof. Notice that in Lemma 3.67 we can safely remove "$\wedge P(x \cdot y)$", because it is automatically implied by $P(x) \wedge P(y)$. A direct inspection then shows that the relevant formulas in point 1 of the previous lemmas are all (equivalent to some formula in) $L_P^{\forall,\mathrm{bdd}}$, and this proves the first part.

In order to prove the second one, let $\psi(\bar{y}) \equiv \forall z \in P\ \varphi(z;\bar{y})$ and suppose that $(N, M) \models \neg\psi(b)$, i.e. for some $a \in M = P(N)$ we have $(N, M) \models \neg\varphi(a;\bar{b})$. Then there is $i < \kappa$ such that $a \in P(N_i) = M_i$ and by L-elementarity of $(M_i)_{i<\kappa}$ for all $i < k$ we have $(N_i, M_i) \models \neg\psi(b)$. On the other hand, if $(N, M) \models \psi(b)$, again by L-elementarity, for all $i < \kappa$ we have $(N_i, M_i) \models \psi(b)$.

The third part follows trivially from the previous two. □

Type-Connected Components of Pairs

Definition 3.72. For (N, M) a saturated pair[21] of models of a NIP theory[22] and $L \subseteq L' \subseteq L_P$, we define $G^{00}_{L'(B)}(N, M)$ to be the intersection of all the bounded index subgroups of M of the form $\Sigma(M; B)$, with Σ a partial L'-type over B, a small tuple from N. Then we define $G^{00}_{L'}(N, M)$ as the intersection of all the former objects as B varies among small tuples of N.

Remark 3.73. The notation $G^{00}_{L'}(N, M)$ may suggest that it is a subgroup of N, with possibly points outside M. Actually it is a subgroup of M *by definition*. The role of N is, morally speaking[23], to provide parameters for external definitions. In order to have a clearer picture, the reader may want to read in advance the statements of Theorem 3.77 and Corollary 3.78 and the proof of the latter.

[21] Recall that, by monster conventions, it is automatically an elementary pair.

[22] We are speaking of M and of N separately. See Remark 3.49.

[23] In the sense that we have not stated this "officially" yet, but our interest is exactly in the case where N is used to build L^{ext}.

To show that $G^{00}_{L'}(N,M)$ exists we need an analogue of Lemma 3.59:

Lemma 3.74. Let (N,M) be a saturated pair, $\Theta(x;\bar{y})$ a partial type and $\bar{b}$ a (possibly infinite) small tuple from N. If $\Theta(M;\bar{b})$ is $\mathrm{Aut}_{L'}(M/\emptyset)$-invariant, then it is definable by a partial L'-type over $\emptyset$.

Proof. We want to write $\Theta(M;\bar{b})$ as the intersection of a family of $L'(\emptyset)$-definable sets with $P(x)$. The natural candidate for such a family is clearly $\Phi = \{\varphi(x) \in L'(\emptyset) \mid \Theta(x;\bar{b}) \wedge P(x) \to \varphi(x)\}$. Of course $\Theta(M;\bar{b}) \subseteq P \cap \bigcap \Phi$; suppose the other inclusion does not hold and, by saturation, let $\tilde{a} \in M$ witness it. If we find $a \in \Theta(M;\bar{b})$ such that $a \equiv_{L'(\emptyset)} \tilde{a}$ then we are done: sending a to $\tilde{a}$ by an element of $\mathrm{Aut}(M/\emptyset)$ contradicts the fact that $\Theta(M;\bar{b})$ should be fixed set-wise by such an automorphism. Suppose there is no such a; then by compactness, there is $\psi(x) \in \mathrm{tp}(\tilde{a}/A)$ such that $\Theta(M;\bar{b}) \subseteq [\neg\psi(x)]$. Then $\neg\psi \in \Phi$, so $\vDash \neg\psi(\tilde{a})$, a contradiction. □

Proposition 3.75 ([CPS14, Proposition 4.13], T NIP)**.** If (N,M) is a saturated pair and $B \subseteq N$ is small, then $G^{00}_{L_P(\emptyset)}(N,M) \subseteq G^{00}_{L(B)}(N,M)$, and as this is true for all small $B \subseteq N$, the same holds for $G^{00}_{L}(N,M)$. In particular the index of the latter is bounded. Moreover, there is a bound on it that only depends on $|L|$.

Proof. Even if $\mathrm{Th}_{L_P}(N,M)$ may not be NIP, formulas of the form $\varphi(x;y,b)\wedge P(x) \wedge P(y)$, with $\varphi(x;y,z) \in L$, are; this is because they correspond to externally definable subsets of M and Corollary 3.57 applies. As we stressed in Remark 3.63, Proposition 3.62 still applies. Taking conjugates of suitable countable subtypes as given by Lemma 3.61 we have that $G^{00}_{L(B)}(N,M)$ is both of bounded index and invariant under L_P-automorphisms. Then, by the previous lemma, it is $L_P(\emptyset)$-type-definable and we have the required inclusion.

As the index of $G^{00}_{L}(N,M)$ is bounded by the one of $G^{00}_{L_P(\emptyset)}(N,M)$, it is sufficient to provide a bound for it. Since $G^{00}_{L_P(\emptyset)}$ is a partial L_P-type over $\emptyset$, it is defined by at most $|L_P| = |L|$ formulas. Suppose that $[M : G^{00}_{L_P(\emptyset)}(N,M)] \geq \lambda$. Let r range among complete L_P-types concentrating on $G^{00}_{L_P(\emptyset)}$ and colour $\{i,j\} \in [\lambda]^2$ with $\Phi_{i,j} \colon r \mapsto \varphi^{i,j}_r$ as in Lemma 3.68. In this case, there are only $|L|$ formulas, hence at most $2^{|L|}$ types r, hence at most $|L|^{2^{|L|}} = 2^{2^{|L|}}$ such Φ. This means that Erdös-Rado is triggered as soon as $\lambda \geq (2^{2^{|L|}})^+$, and yields $[M : G^{00}_{L_P(\emptyset)}(N,M)] \leq 2^{2^{|L|}}$. □

Corollary 3.76. For some small $B \subseteq N$ we have $G^{00}_{L}(N,M) = G^{00}_{L(B)}(N,M)$.

Proof. Otherwise we would have too many subgroups H such that $G \supseteq H \supseteq G^{00}_{L}(N,M)$, contradicting fact that the latter has bounded index. □

We are finally ready for the crucial result.

Theorem 3.77 ([CPS14, Theorem 4.14], T NIP). Let (N, M) be a saturated pair. Then $G^{00}(M) = G^{00}_L(N, M)$.

Before proving it, let us see its consequences.

Corollary 3.78 ([CPS14, Corollary 4.16], T NIP). Let $\widetilde{\mathfrak{U}}$ be a monster model for $\mathrm{Th}(M^{\mathrm{ext}})$, and let $\mathfrak{U} = \widetilde{\mathfrak{U}} \upharpoonright L$ be its reduct to the original language. Then $G^{00}(\widetilde{\mathfrak{U}}) = G^{00}(\mathfrak{U})$.

Proof. Use some $N \ {}^{+}\!\succ M$ to define L^{ext} and let $(N', \widetilde{\mathfrak{U}})$ be a monster for (N, M). Then $G^{00}_L(N', \mathfrak{U})$ is exactly $G^{00}(\widetilde{\mathfrak{U}})$, and by Theorem 3.77 it equals $G^{00}(\mathfrak{U})$. □

Remark 3.79. Again by the fact that two compact Hausdorff topologies are either incomparable or identical, the logic topology on G/G^{00} evaluated with respect to L is the same as the logic topology evaluated with respect to L^{ext}.

Proof of Theorem 3.77. The strategy of proof is to fix some $\lambda \ggg [M : G^{00}_L(N, M)]$ and set up an iterative construction that, if carried out for $\lambda + 1$ steps, contradicts Proposition 3.75[24]. When this happens, we will have an M_α that will allow us to prove the inclusion $G^{00}(M) \subseteq G^{00}_L(N, M)$. The other one is trivial.

The construction goes as follows. For all stages $\alpha \leq \lambda$, find models M_α, N_α, a partial L-type $\Sigma_\alpha(x; \bar{y}_\alpha)$ and a $|\bar{y}_\alpha|$-tuple $\bar{b}_\alpha \in N_\alpha$ with the following properties:

1. $M_0 = M$, $N_0 = N$, and $\Sigma_0(M_0; \bar{b}_0) = G^{00}_L(N, M)$.
2. $(M_\beta \mid \beta \leq \alpha)$ and $(N_\beta \mid \beta \leq \alpha)$ are two L-elementary chains.
3. (N_α, M_α) is a sufficiently saturated elementary pair.
4. $\Sigma_\alpha(x; \bar{y}_\alpha)$ is a partial L-type of small size with respect to (N_α, M_α).
5. For all $\gamma \geq \alpha$, $\mathrm{tp}((\bar{b}_i)_{i \leq \alpha} / L_P^{\forall,\mathrm{bdd}})$ is the same in M_γ as in M_α.
6. $\Sigma_\alpha(x; \bar{b}_\alpha)$ is an hereditary subgroup of $P(x)$ of hereditarily bounded index in (N_α, M_α).
7. For all $j < i \leq \alpha$ we have the strict inclusion $\Sigma_i(M_\alpha; \bar{b}_i) \subsetneq \Sigma_j(M_\alpha; \bar{b}_j)$.

As anticipated, if this construction can be done all the way up to stage λ, then in (N_λ, M_λ) we have the chain of subgroups $(\Sigma(M_\lambda; \bar{b}_i) \mid i < \lambda)$, which are of bounded index and defined by small partial types. Therefore they all include $G^{00}_L(N_\lambda, M_\lambda)$, contradicting Proposition 3.75 as anticipated[25].

[24] This is where the NIP stops the construction.
[25] Here we use the fact that the bound does not depend on N, M.

Claim. The first stage α^* that cannot be completed is a successor.

Proof of the Claim. Limit stages can always be carried out in the following way: if α^* is a limit, take as (N^*_α, M^*_α) any sufficiently saturated extension of the elementary pair $(\bigcup_{\alpha<\alpha^*} N_\alpha, \bigcup_{\alpha<\alpha^*} M_\alpha)$ and then set $\bar{b}_{\alpha^*} = \bigcup_{\alpha<\alpha^*} \bar{b}_\alpha$ and $\Sigma_{\alpha^*}(x; \bar{b}_{\alpha*}) = \bigcup_{\alpha<\alpha^*} \Sigma_\alpha(x; \bar{b}_\alpha)$. Then, by inductive hypothesis and Lemma 3.71, these objects still satisfy the requirements of the construction. $\square_{\text{CLAIM}}$

Write $\alpha^* = \alpha + 1$ and let $K \succ N_\alpha \succ M_\alpha$ be very saturated. Then define $\mathcal{F}$ as the family of partial types in $L(K)$ defining in (K, M_α) an hereditary subgroup of M_α of hereditarily bounded index, and let[26] $\Lambda(x) = \bigcup_{\Sigma(x)\in\mathcal{F}} \Sigma(x)$. We then have

$$\Lambda(M_\alpha) \subseteq G^{00}_L(N_\alpha, M_\alpha) \subseteq \Sigma_\alpha(M_\alpha; \bar{b}_\alpha)$$

where the first inclusion is trivial by definition of Λ and the second one is again Lemma 3.71, which can be applied since $\mathrm{tp}((\bar{b}_i)_{i\le\alpha})/L_P^{\forall,\mathrm{bdd}})$, only caring about what happens in M_α, is the same in (K, M_α) as in (N_α, M_α).

Claim. $\Lambda(M_\alpha) = G^{00}_L(N_\alpha, M_\alpha)$. Moreover, it is $\mathrm{Aut}_L(M_\alpha/\emptyset)$-invariant.

Proof of the Claim. We show that if the inclusion is strict, then the construction could have continued after α. Take as $(N_{\alpha+1}, M_{\alpha+1})$ a sufficiently saturated elementary extension of (K, M_α). Since $\Lambda(x)$ is defined as the intersection of at most $2^{|M|_\alpha}$ hereditary subgroups of hereditarily bounded index, it is small with respect to $(N_{\alpha+1}, M_{\alpha+1})$; we then take as $\bar{b}_{\alpha+1}$ the parameters involved in defining $\Lambda(x)$ and write it as $\Sigma_{\alpha+1}(x; \bar{b}_{\alpha+1})$. Again by Lemma 3.71, these new objects still satisfy the hypotheses of the construction.

Now let $f \in \mathrm{Aut}_L(M_\alpha/\emptyset)$. Since K is saturated[27] enough, we can extend it to some $\tilde{f} \in \mathrm{Aut}_{L_P}(K, M_\alpha/\emptyset)$. But now $\tilde{f}(\Lambda(M_\alpha)) = \Lambda(M_\alpha)$ by definition of $\Lambda(x)$. $\square_{\text{CLAIM}}$

By the latest Claim and Lemma 3.74 $G^{00}_L(N_\alpha, M_\alpha)$ is L-type-definable over M_α. We can therefore start the following chain of inclusions:

$$G^{00}(M_\alpha) \subseteq G^{00}_L(N_\alpha, M_\alpha) \subseteq \Sigma_\alpha(M_\alpha; \bar{b}_\alpha) \subseteq \Sigma_0(M_\alpha; \bar{b}_0)$$

Since $G^{00}(M_\alpha)$ is definable over $\emptyset$ by Theorem 3.64, the inclusion $G^{00}(-) \subseteq \Sigma_0(-; \bar{b}_0)$ is still valid with M_0 in place of M_α. Since $\Sigma_0(M_0; \bar{b}_0) = G^{00}_L(N, M)$,

[26] The union is intended at the level of formulas; contravariantly, it corresponds to the intersection of the defined subgroups.

[27] Recall that we stipulated in Notation 3.65 that "saturated" means "saturated and strongly homogeneous".

the proof is completed, but there is more: since, unless $\alpha = 0$, by requirement 7 the last inclusion in the chain above should be strict, we have that $\alpha^* = 1$ and the construction cannot even start, otherwise $G^{00}(M_0) \subseteq \Sigma_0(M_0; \bar{b}_0) = G^{00}_L(N, M)$ would be strict too. $\square$

3.6 Invariant Heirs

The aim of this section is to show that in NIP theories types in $S(M)$ have global heirs which are M-non-forking, hence M-invariant by NIP. This will follow from a special case of the so-called *Broom Lemma*, which was proven in [CK12] for the broader class of NTP_2 theories. Since we only need it in the NIP case we follow the proof of [Sim15] and, in order to simplify some details, do not work in full generality.

The Tree Property of The Second Kind

The definition of NTP_2 is in the same spirit of the definition of NIP: we look at a partitioned formula as a uniformly definable family, and we isolate a "wild" combinatorial property of this family; then we ask the formula to omit it in order to be "tame".

Definition 3.80. A partitioned formula $\varphi(x; y)$ has the *tree property of the second kind*, or TP_2, iff there is $k \in \omega$ and an infinite matrix of $|y|$-tuples $(b_i^t \mid i < \omega, t < \omega)$ such that for all the rows t the partial type $\{\varphi(x; b_i^t) \mid i < \omega)\}$ is k-inconsistent, but every time we choose an element from each row with some $\eta\colon \omega \to \omega$, the partial type $\{\varphi(x; b^t_{\eta(t)}) \mid t < \omega\}$ is consistent. A formula has/is NTP_2 iff it does not have TP_2. A theory is NTP_2 iff all of its formulas are.

Fact 3.81. If $\varphi(x; y)$ has TP_2 with $k = 2$, then it has IP.

Proof. For each $X \in \mathscr{P}(\omega)$, let η_X be its characteristic function and let a_X realize $\{\varphi(x; b^t_{\eta(t)}) \mid t < \omega\}$, which is consistent by hypothesis. In other words, concentrate the attention on the η with values on the first two columns. Then, by 2-inconsistency of the rows, for any t the only b_i^t satisfying $\varphi(a_X; y)$ is the one with $i = \eta(t)$. Therefore, looking at the second column $i = 1$, we discover that $\vDash \varphi(a_X; b_1^t) \iff \eta_X(t) = 1 \iff t \in X$. $\square$

It can be proven (see for example [She90, Chapter 4, § III, Theorem 7.7]) that in order to show that a theory is NTP_2 it suffices to show the failure of TP_2 in the special case above, i.e. where $k = 2$. We will not use Fact 3.81; instead we will arrive at the result that NIP theories are NTP_2 via a small detour in the realm of *mutually indiscernible sequences*.

Mutually Indiscernible Sequences

Definition 3.82. The sequences[28] in a family $(I_t \mid t \in X)$ are called *mutually indiscernible over* A iff each I_t is $A \cup \{I_s \mid s \in X \setminus \{t\}\}$-indiscernible.

We now want to prove that in NIP theories, for any $b \in \mathfrak{U}$, a sufficiently large family of A-mutually indiscernible sequences always has a subfamily which is mutually indiscernible over Ab, because this will be shown to fail in a TP_2 theory. This will require some preliminary results. The first one is that, in the same fashion as we "turn" any sequence to an indiscernible one with the Standard Lemma, it is possible to build families of mutually indiscernible sequences with prescribed Ehrenfeucht-Mostowski behaviour.

Construction 3.83. Let $(I_t \mid t < \alpha)$ be a family of sequences. Inductively define $A_t = A \cup \{I'_s \mid s < t\} \cup \{I_s \mid s > t\}$ and, by the Standard Lemma, replace I_t with an A_t-indiscernible $I'_t \vDash \mathrm{EM}(I_t/A_t)$.

Lemma 3.84 ([Sim15, Lemma 4.2])**.** The sequences $(I'_t \mid t < \alpha)$ of Construction 3.83 are mutually indiscernible over A.

Proof. The only thing that could go wrong is that a stage messes up with the previous ones. In other words, we have to make sure that the sequence I'_s built at stage s will remain indiscernible even with respect to the I'_t that we will build later for $t > s$. We claim that this automatically happens.

If not, let t be the first stage for which there is some $s < t$ such that I'_s is not indiscernible with respect to a formula $\varphi(x)$ with parameters in $A \cup \{I'_r \mid r < t, r \neq s\} \cup I'_t$. Make the parameters from I'_t explicit, i.e. write $\varphi(x) = \varphi(x; b)$ where $b \in I'_t$ and $\varphi(x; y)$ has hidden parameters from $A \cup \{I'_r \mid r < t, r \neq s\}$. Let $\bar\imath = i_0 < \ldots < i_{n-1}$ and $\bar\jmath = j_0 < \ldots < j_{n-1}$ be such that $\vDash \psi(b) = \varphi(a_{\bar\imath}; b) \triangle \varphi(a_{\bar\jmath}; b)$, with $a_{\bar\imath}, a_{\bar\jmath} \in I'_s$. Notice that by inductive hypothesis b cannot be the empty tuple, and that it can be replaced with an element of I_t: if for all $\tilde b \in I_t$ we had $\vDash \neg\psi(\tilde b)$, the same would be true for b, because then $\neg\psi(y) \in \mathrm{EM}(I_t/A \cup \{I'_r \mid r < t\}) \subseteq \mathrm{EM}(I'_t/A \cup \{I'_r \mid r < t\})$. Inductively, we can replace all parameters in $\varphi(x_0) \triangle \varphi(x_1)$ living in some I'_r with parameters in I_r by exploiting the fact that $I'_r \vDash \mathrm{EM}(I_r/A_r)$. After at most $\min\{|t - s|, |b|\}$ steps this will violate A_s-indiscernibility of I'_s. $\square$

We first need the following consequence of the existence of honest definitions.

Lemma 3.85 ([Sim15, Corollary 3.18], T NIP)**.** If $P \subseteq M$, $b \in M$ is a finite tuple and $(M, P) \prec^+ (M', P')$, there is $P_0 \subseteq P'$ of size at most $|T|$ such that for all $a_0, a_1 \in P$ we have $a_0 \equiv_{P_0} a_1 \Rightarrow a_0 \equiv_b a_1$.

[28] Not necessarily of the same length.

Proof. For all $\varphi(x;y) \in L$ choose an honest definition $\psi_\varphi(x;d_\varphi)$ of $\varphi(x;b)$ inside (M',P'). This implies that for all $a \in A$ we have $\vDash \varphi(a;b) \leftrightarrow \psi_\varphi(a;d_\varphi)$. Let $P_0 = \{d_{\varphi,b} \mid \varphi(x;y) \in L\}$ and notice that it has the required cardinality. □

Proposition 3.86 ([Sim15, Proposition 4.8], T NIP)**.** Let $(I_t \mid t \in X)$ be family of sequences of n-tuples of M, mutually indiscernible over A, and let $b \in M$ be a finite tuple. There is $X_b \subseteq X$ such that $|X_b| \leq |T|$ and the $(I_t \mid t \in X \setminus X_b)$ are mutually indiscernible sequences over Ab.

Proof. Up to passing to M^{eq} and using Theorem 3.32 we can assume that $n = 1$, i.e. that each $I_t = (a_i^t \mid i \in J_t)$ is a sequence of single elements. Let $L' = L \cup \{A(x), P(x), R(x,y)\}$ and expand M to an L'-structure with the following interpretations

$$A(M) = A \qquad P(M) = A \cup \{a_i^t \mid t \in X, i \in J_t\}$$
$$R(M^2) = \{(a_i^t, a_j^t) \mid t \in X, i, j \in J_t, i < j\}$$

Take some $(M', A', P', R') \mathrel{^{+}\succ} (M, A, P, R)$. By L'-elementarity, there are $X' \supseteq X$ and for each $t \in X'$ some J'_t such that

- $P' = P(M') = A' \cup \{a_i^t \mid t \in X', i \in J'_t\}$;
- if $t \in X$ then $J_t \subseteq J'_t$;
- if $I'_t = (a_i^t \mid i \in J'_t)$ then the family $(I'_t \mid t \in X')$ is made of mutually indiscernible sequences over A'.

Use Lemma 3.85 to find $P_0 \subseteq P'$ such that $|P_0| \leq |T|$ and $a_0 \equiv_{P_0} a_1 \Rightarrow a_0 \equiv_b a_1$ holds for all $a_0, a_1 \in P(M)$, and discard the $t \in X$ involved in P_0, i.e. let $X_b = \{t \in X \mid \exists i \in J_t\ a_i^t \in P_0\}$. Notice that $|X_b| \leq |P_0| \leq |T|$.

To conclude, let $\varphi(x; y_0, y_1, y_2) \in L$, $a \in A$, $t_0 \in X \setminus X_b$, d a tuple from $\{a_i^t \mid t \in X \setminus (X_b \cup t_0), i \in J_t\}$ and fix $a_{\bar\imath}^{t_0}$, $a_{\bar\jmath}^{t_0}$ for $\bar\imath = i_0 < \ldots < i_n$ and $\bar\jmath = j_0 < \ldots < j_n$. The hypotheses on $(I_t \mid t \in X)$ imply that $(a, a_{\bar\imath}^{t_0}, d)$ and $(a, a_{\bar\jmath}^{t_0}, d)$ have the same type over P_0. Therefore, since they live in P, they have the same type over b and this proves $\vDash \varphi(a_{\bar\imath}^{t_0}; a, d, b) \leftrightarrow \varphi(a_{\bar\jmath}^{t_0}; a, d, b)$. □

We invite the reader to notice that results such as Proposition 3.86, Corollary 3.58, Theorem 3.34, or Corollary 3.22 all have a similar flavour: in NIP theories, up to discarding a small amount (less than $|T|^+$ sequences, an initial segment…) of objects, some constructions carry more information than expected ("more" indiscernibility, "double externals" already externals, strong Borel-definability, etc.). Also, compare Proposition 3.86 with Lemma 3.29.

The Broom Lemma for NIP Theories

Proposition 3.87 ([Sim15, Proposition 5.31])**.** NIP theories are NTP$_2$.

Proof. Let T have TP$_2$. Since Ehrenfeucht-Mostowski types are aware of k-inconsistency, by Lemma 3.84 and compactness we can find a matrix made of mutually $\emptyset$-indiscernible sequences that witnesses TP$_2$ apart from having $|T|^+$ rows instead of ω. Let b realize the partial type corresponding to the first column, namely $\{\varphi(x; b_0^t) \mid t < |T|^+\}$. Then no row t can be b-indiscernible, because this would contradict the fact that $\{\varphi(x; b_i^t) \mid i < \omega\}$ is k-inconsistent. Therefore Proposition 3.86 fails for b, and T cannot be NIP. □

In order to explore the consequences of the previous fact, we need a preliminary result.

Lemma 3.88 ([Sim15, Lemma 5.34], T any theory)**.** If $\bar{b} = (b_i \mid i < \omega)$ is A-indiscernible and $p = \mathrm{tp}(\bar{b}/A)$ does not fork over A, then there is a global extension of p that is the type of some $\mathfrak{U}$-indiscernible sequence and does not fork over A.

Proof. Let $\bar{c} = (c_i \mid i < \omega)$ realize an A-non-forking global extension of p. Clearly $\bar{c}$ need not be $\mathfrak{U}$-indiscernible, but it suffices to use the Standard Lemma to find a $\mathfrak{U}$-indiscernible $\bar{d} \models \mathrm{EM}(\bar{c}/\mathfrak{U}) \supseteq \mathrm{EM}(\bar{c}/A) = p$. To show that $q = \mathrm{tp}(\bar{d}/\mathfrak{U})$ does not fork over A, notice that if $q \vdash \psi(\vec{x})$ then there is $\vec{c} \subseteq \bar{c}$ such that $\psi(\vec{c})$, otherwise we would have had $q \vdash \neg\psi$ by definition of Ehrenfeucht-Mostowski type. Therefore ψ cannot fork over A. □

Lemma 3.89 ([Sim15, Lemma 5.35], T NIP)**.** Dividing over extension bases can be witnessed by Morley sequences. More precisely, let A be an extension base and let $\varphi(x; b)$ divide over A. Then there are $M \supseteq A$, $p \in S^{\mathrm{inv}}(\mathfrak{U}, M)$ and a Morley sequence $\bar{c}$ of p over M such that $c_0 = b$ and $\{\varphi(x; c_i) \mid i < \omega\}$ is inconsistent.

Proof. Let $\bar{b}$ witness A-dividing. Since A is an extension base, $\mathrm{tp}(\bar{b}/A)$ does not fork over A, and by the previous lemma it has a global A-non-forking extension q which is the type of a $\mathfrak{U}$-indiscernible sequence. Since q does not fork over A, it is M-invariant for any $M \supseteq A$ by Corollary 3.37. Let $(\bar{c}^t \mid t < \omega)$ be a Morley sequence of q over M. For each $t < \omega$, since $\bar{c}^t \equiv_A \bar{b}$, the partial type $\{\varphi(x; c_i^t) \mid i < \omega\}$ is k-inconsistent. Then, by Proposition 3.87, there is $\eta\colon \omega \to \omega$ such that $\{\varphi(x; c_{\eta(t)}^t) \mid t < \omega\}$ is inconsistent, and since q was the type of a $\mathfrak{U}$-indiscernible sequence, this is a Morley sequence of $\mathrm{tp}(c_0^0/\mathfrak{U})$ over M. Up to an element of $\mathrm{Aut}(\mathfrak{U}/A)$ (that may change M, but this is not a problem) we have $c_0 = b$. □

What follows is, as anticipated, a special case of a result called the *Broom Lemma.* For the general statement see [CK12, Lemma 3.7].

Lemma 3.90 ([Sim15, Lemma 5.36 (Broom Lemma)], T NIP)**.** Let A be an extension base and $\pi(x)$ be a partial Lstp_A-invariant[29] type. Let $\psi(x;b)$ and $\{\varphi_i(x;c) \mid i<n\}$ be such that each φ_i divides over A and $b \mathop{\smile\!\!\!\!\!\!\mid}_A c$, and suppose that

$$\pi(x) \vdash \psi(x;b) \vee \bigvee_{i<n} \varphi_i(x;c)$$

Then $\pi(x) \vdash \psi(x;b)$.

Proof. By induction on n. If $n=0$ there is nothing to prove, so assume that the Broom Lemma is true for n. By Lemma 3.89 let $(c_j)_{j<\omega}$ be a Morley sequence of some $p \in S^{\mathrm{inv}}(\mathfrak{U}, M)$ over $M \supseteq A$ witnessing dividing of φ_n, say with respect to k. By Lemma 2.85, modulo an $f \in \mathrm{Autf}(\mathfrak{U}/A)$, we can assume that $b \mathop{\smile\!\!\!\!\!\!\mid}_A (c_j)_{j<\omega}$. This will prove that $f(\pi(x)) \vdash \psi(x; f(b))$, but since π is Lstp_A-invariant it is enough. Since now $\mathrm{tp}(b/A(c_j)_{j<\omega})$ does not fork over A, it is Lstp_A-invariant, hence $(c_j)_{j<\omega}$ is Ab-indiscernible[30]. This implies that $(bc_j)_{j<\omega}$ is A-indiscernible, and again by Lstp_A-invariance of $\pi(x)$, for all $j<\omega$ we have

$$\pi(x) \vdash \psi(x;b) \vee \bigvee_{i<n+1} \varphi_i(x;c_j)$$

Taking the conjunction over the first k among the $(c_j)_{j<\omega}$ we have

$$\pi(x) \vdash \psi(x;b) \vee \bigwedge_{j<k} \bigvee_{i<n+1} \varphi_i(x;c_j)$$

Since the $(c_j)_{j<\omega}$ were chosen specifically in order for $\{\varphi_n(x;c_j) \mid j<\omega\}$ to be k-inconsistent[31], no realization of $\pi(x)$ can satisfy the conjunction by always realizing the last disjoint. In other words,

$$\pi(x) \vdash \psi(x;b) \vee \bigvee_{\substack{j<k \\ i<n}} \varphi_i(x;c_j) \tag{3.1}$$

Recall that we arranged to have $b \mathop{\smile\!\!\!\!\!\!\mid}_A (c_j)_{j<\omega}$. For all $j<k$ this implies[32] $b \mathop{\smile\!\!\!\!\!\!\mid}_{Ac_{>j}} c_j$, and since $(c_j)_{j<\omega}$ is a Morley sequence of p over A by Lemma 2.84 we also have $c_{>j} \mathop{\smile\!\!\!\!\!\!\mid}_A c_j$. Therefore $bc_{>j} \mathop{\smile\!\!\!\!\!\!\mid}_A c_j$ by Lemma 2.82. For $j=0$ this allows us to apply the induction hypothesis by rewriting (3.1) as

$$\pi(x) \vdash \Big(\psi(x;b) \vee \bigvee_{\substack{0<j<k \\ i<n}} \varphi_i(x;c_j)\Big) \vee \bigvee_{i<n} \varphi_i(x;c_0)$$

[29] See Remark 2.51.

[30] Since for $\bar{\imath} < \bar{\jmath}$ we have that $a_{\bar{\imath}}a_{\bar{\jmath}}$ starts an A-indiscernible sequence, by Lstp_A-invariance for all $\varphi(x;y) \in L(A)$ we have $\vDash \varphi(b;a_{\bar{\imath}}) \leftrightarrow \varphi(b;a_{\bar{\jmath}})$.

[31] ... and since this is preserved by any $f \in \mathrm{Autf}(\mathfrak{U}/A)$...

[32] Trivially, if $\mathrm{tp}(b/Ac_{\geq j})$ forks over $Ac_{>j}$, it also forks over A.

and we can "sweep away" the last disjunction. We can then rewrite the previous formula again, set $j = 1$ and iterate; after $k - 1$ more "sweeps", we end up with $\pi(x) \vdash \psi(x)$. □

Corollary 3.91 ([Sim15, Corollary 5.37], T NIP)**.** Let A be an extension base and $\pi(x)$ be a partial Lstp_A-invariant type. Let $\psi(x; b)$ and $\{\varphi_i(x; c) \mid i < n\}$ be such that each φ_i divides over A and, and suppose that

$$\pi(x) \vdash \psi(x; b) \vee \bigvee_{i<n} \varphi_i(x; c)$$

Then there are $(b_j \mid j < m)$ such that for all $j < m$ we have $b_j \equiv_{\mathrm{Lstp}_A} b$ and

$$\pi(x) \vdash \bigvee_{j<m} \psi(x; b_j)$$

Proof. Apart from the fact that this time we do not know $b \mathop{\smile\!\!\!\!\!\!\mid}_A c$, our hypotheses are the same as before. The type $\tilde{\pi}(x) = \pi(x) \cup \{\neg\psi(x; \tilde{b}) \mid \tilde{b} \equiv_{\mathrm{Lstp}_A} b\}$ is still Lstp_A-invariant, and trivially $\emptyset \mathop{\smile\!\!\!\!\!\!\mid}_A c$, so we can apply the Broom Lemma to $\tilde{\pi}(x)$ and $x \neq x \vee \bigvee_{i<n} \varphi_i(x; c)$, and discover that $\tilde{\pi}(x) \vdash x \neq x$. Now apply compactness. □

Global Invariant Heirs

We are ready to reap the fruits of this section.

Theorem 3.92 ([Sim15, Proposition 5.39], T NIP)**.** Every type in $S(M)$ has a global, M-non-forking heir.

Proof. Let $p \in S_x(M)$ and consider the following partial types

$$\pi_0(x) = \{\psi(x; b) \in L(\mathfrak{U}) \mid \forall \tilde{b} \in M \; p \vdash \psi(x; \tilde{b})\}$$
$$\pi_1(x) = \{\varphi(x; c) \in L(\mathfrak{U}) \mid \neg\varphi(x; c) \text{ divides over } M\}$$

up to rewriting $\pi_0(x)$ as

$$\pi_0(x) = \{\psi(x; b) \in L(\mathfrak{U}) \mid \neg\exists \tilde{b} \in M \; p \vdash \neg\psi(x; \tilde{b})\}$$

it is clear that all we have to show is that $p(x) \cup \pi_0(x) \cup \pi_1(x)$ is consistent. If it is not, since π_0 is closed under conjunctions, by compactness we can find $\psi(x; b) \in \pi_0(x)$ and $\{\varphi_i(x; c) \mid i < m\} \subseteq \pi_1(x)$ such that

$$p(x) \vdash \neg\psi(x; b) \vee \bigvee_{i<n} \neg\varphi_i(x; c)$$

Since models are extension bases by Proposition 2.87 and p, as an element of $S_x(M)$, is M-invariant, we are in the hypotheses of the previous Corollary and there are, for $j < m$, some $b_j \equiv_M b$ such that $p(x) \vdash \bigvee_{j<m} \neg\psi(x; b_j)$.

Suppose that we find a formula $\theta(y) \in \operatorname{tp}(b_0, \dots, b_{m-1}/M)$ such that $p(x) \vdash \forall y \; (\theta(y) \to \bigvee_{j<m} \neg\psi(x; y_j))$. Then, since M is a model, there are some $\tilde{b}_0, \dots, \tilde{b}_{m-1} \in M$ realizing $\theta(y)$, and this contradicts $\psi(x; b) \in \pi_0(x)$. Suppose that for all formulas $\theta(y) \in \operatorname{tp}(b_0, \dots, b_{m-1}/M)$ we have

$$p(x) \vdash \exists y \left(\theta(y) \wedge \neg \bigvee_{j<m} \neg\psi(x; y_j)\right)$$

Then by compactness[33] there is $(d_0, \dots, d_{m-1}) \equiv_M (b_0, \dots, b_{m-1})$ such that $p(x) \vdash \neg \bigvee_{j<m} \neg\psi(x; d_j)$. If $f \in \operatorname{Aut}(\mathfrak{U}/M)$ is such that $f(d_j) = b_j$, we have $p = f(p) \vdash \neg \bigvee_{j<m} \neg\psi(x; b_j)$, a contradiction. □

Corollary 3.93 (T NIP)**.** Every type in $S(M)$ has a global, M-invariant heir.

Proof. By the previous proposition and Theorem 3.38. □

For the reader's convenience, we state the following result, which can be found in [CK12, Theorem 3.9] and will not be used in this thesis. A proof for the special case of NIP theories can be found in [Sim15, Theorem 5.49], and it requires *strict non-forking*.

Theorem 3.94. In NTP$_2$ theories, forking over an extension base equals dividing. In particular this is true over models.

[33] It suffices to show that the family of closet sets of $S_{xy}(M)$ of the form $[p(x)] \cap [\theta(y)]$ has non-empty intersection, find a point inside it, and then take the restriction to $S_y(M)$. But we just showed that this family has the FIP.

Chapter 4

Definable Amenability

A group G is *amenable* if there is a finitely additive probability measure μ defined on the whole $\mathscr{P}(G)$ which is left-translation-invariant, i.e. for all $g \in G$ and $X \subseteq G$ we have $\mu(X) = \mu(g \cdot X)$. The notion was introduced by von Neumann in his work on the Banach-Tarski paradox, which is essentially due to the fact that $\mathrm{SO}(3, \mathbb{R})$ is not amenable, as it contains a copy of the free group on two generators F_2. This allows to decompose the three-dimensional ball in a finite number of pieces $D^3 = \bigsqcup_{i<k} X_i$ that can be then "reassembled" via rigid motions in two copies of the same ball, say $D^3 \sqcup (D^3 + (4, 0, 0)) = \bigsqcup_{i<k} g_i X_i$, where the g_i are suitable orientation-preserving isometries of $\mathbb{R}^3$. Notice that at least some of the X_i must be non-Lebesgue-measurable, since such a decomposition would violate the fact that the Lebesgue measure is translation-invariant. This is a key issue, since as soon as one is only concerned with definable sets, these kind of obstructions to amenability disappear, and both $\mathrm{SO}(3, \mathbb{R})$ and F_2 are *definably* amenable.

Amenable groups were extensively studied in the past century and continue to be an active research topic; as a result, there is a vast literature on the subject. Some classical results are that all abelian, and more generally all solvable groups are amenable. On the other hand, a group containing a copy of F_2 is not amenable[1]. A systematic study of amenability is beyond the scope and purpose of this thesis; the interested reader can consult for instance the monograph [Pie84].

Permanent Assumption 4.1. From now on, G will be an $\emptyset$-definable group.

4.1 Definably Amenable Groups

Notation 4.2. Let A be a set of parameters and identify G with a clopen subset of $S_x(A)$ for a suitable x. We denote with $\mathfrak{M}_G(A)$ the closed sub-

[1] The converse of this statement was known under the name of *von Neumann's conjecture*, and it was disproven.

space of $\mathfrak{M}_x(A)$ consisting of measures concentrating on G, i.e. such that the formula defining G has measure one.

Equivalently, this means $\mu \in \mathfrak{M}_G(A) \leftrightarrow S(\mu) \subseteq G$. See also Lemma 2.102.

Definition 4.3. G is *definably amenable* iff it has a global left-invariant measure, i.e. iff there is some $\mu \in \mathfrak{M}_G(\mathfrak{U})$ such that for all $g \in \mathfrak{U}$ and Borel $B \subseteq S_G(\mathfrak{U})$ we have $\mu(g \cdot B) = \mu(B)$. We also say that μ *is* $G(\mathfrak{U})$*-invariant.*

So the difference between "amenable" and "definably amenable" resides in asking the invariant measure to be defined *only on definable subsets of* G.

Remark 4.4. The focus on *left* invariant measures is convenient, but not essential. In fact, if $\mu(x)$ is left-invariant, then $\nu(\varphi(x)) = \mu(\varphi(x^{-1}))$ is obviously right-invariant. Anyway this does *not* mean that all left-invariant measures are also right-invariant; see [CSed, Example 6.1].

Lemma 4.5. Definable amenability may be checked on a small model. More precisely every $G(M)$-invariant $\mu \in \mathfrak{M}_G(M)$ can be extended to a $G(\mathfrak{U})$-invariant $\nu \in \mathfrak{M}_G(\mathfrak{U})$. Conversely, the restriction of a $G(\mathfrak{U})$-invariant $\nu \in \mathfrak{M}_G(\mathfrak{U})$ to $\mathfrak{M}_G(M)$ is $G(M)$-invariant.

Proof. Use Fact 2.108 to lift such a μ to a measure on a sufficiently saturated $\widetilde{\mathfrak{U}}$, an notice that this lifting will be $G(\widetilde{\mathfrak{U}})$-invariant by $L(\widetilde{M_\mu})$-elementarity with respect to the formula $\forall g \in G\ f_{\varphi(y^{-1}\cdot x)}(e) = f_{\varphi(y^{-1}\cdot x)}(g)$. Then embed $\mathfrak{U}$ into $\widetilde{\mathfrak{U}} \restriction L$ and take a restriction. The second statement is trivial. □

Clearly, all amenable groups are automatically definably amenable regardless of the first-order structure we put on them, since we already have a left-invariant (finitely additive) measure defined on their whole power-set, and as a special case on all the definable subsets. As anticipated, F_2 and $\mathrm{SO}(3,\mathbb{R})$ are examples of definably amenable groups that are not amenable. This is due to two more general facts.

The first one is that stable groups are definably amenable (see [Sim15, Example 8.14]) and F_2 happens to be stable (see [Sel13]). The second one is that $\mathrm{SO}(3,\mathbb{R})$ is a definably compact o-minimal group and, in an o-minimal structure, a group G is definably amenable if and only if (see [CP12]) there are a torsion free H, a definably compact C, and a short exact sequence $1 \to H \to G \to C \to 1$. In this case it suffices to take $\mathrm{SO}(3,\mathbb{R}) = G = C$ and $H = 1$.

An example of a group which is not even definably amenable is $\mathrm{SL}(2,\mathbb{R})$. This follows from the previous characterization, as it is well-known that its only normal subgroup is its center $\{I, -I\}$ and the quotient $\mathrm{PSL}(2,\mathbb{R})$ is not compact. As anticipated, in [GPP15] it was proven that the Ellis group of $\mathrm{SL}(2,\mathbb{R})$ is *not* isomorphic to $\mathrm{SL}(2,\mathbb{R})/\mathrm{SL}(2,\mathbb{R})^{00}$, hence the "definably amenable" hypothesis in the Ellis Group Conjecture is necessary.

For more about definably amenable groups in NIP theories, see [CSed].

4.2 Extension of Measures to Shelah's Expansion

Permanent Assumption 4.6. From now until the end of the chapter the ambient theory will be supposed NIP.

Our first concern is to prove that definably amenability is preserved after passing to M^{ext}. This amounts to extending a $G(M)$-invariant $\mu \in \mathfrak{M}_G(M)$ to a $G(M)$-invariant $\tilde{\mu} \in \mathfrak{M}_G(M^{\mathrm{ext}})$, and was done in [CPS14]. The strategy of proof is the following:

1. Prove that every $G(M)$-invariant measure in $\mathfrak{M}_G(M)$ has a $G(\mathfrak{U})$-invariant extension in $\mathfrak{M}_G^{\mathrm{inv}}(\mathfrak{U}, M)$. As a special case, it will be $G(M)$-invariant.
2. Take a retraction of $\mathfrak{M}^{\mathrm{inv}}(\mathfrak{U}, M)$ onto $\mathfrak{M}^{\mathrm{fs}}(\mathfrak{U}, M)$ that preserves the property of being $G(M)$-invariant.
3. Identify $\mathfrak{M}^{\mathrm{fs}}(\mathfrak{U}, M)$ with $\mathfrak{M}(M^{\mathrm{ext}})$.

Heirs of a Measure

A $G(\mathfrak{U})$-invariant, M-invariant extension will be found with the help of heirs. A notion of *heir of a measure* was given in [HPP08], but we need a stronger definition[2].

Definition 4.7 ([CPS14, Definition 2.2])**.** Let $\mu \in \mathfrak{M}_x(M)$ and $\nu \in \mathfrak{M}_x(\mathfrak{U})$. We say that ν is an *heir* of μ iff for all $a \in \mathfrak{U}$, $n \in \omega$ and, for $i < n$, $\varphi_i(x; y) \in L(M)$ and $r_i \in [0, 1)$ there is $b \in M$ such that

$$\bigwedge_{i<n} \nu(\varphi_i(x; a)) > r_i \implies \bigwedge_{i<n} \mu(\varphi_i(x; b)) > r_i$$

Clearly, if μ is a type, this coincides with Definition 1.71. This is what we will use heirs for:

Proposition 4.8 ([CPS14, Proposition 3.4])**.** Global heirs of $G(M)$-invariant measures over M are $G(\mathfrak{U})$-invariant.

Proof. Being an heir allows to produce, from a counterexample to $G(\mathfrak{U})$-invariance, a counterexample to $G(M)$-invariance. Let us spell this out. Let $\mu \in \mathfrak{M}_x(M)$ be $G(M)$-invariant and let $\nu \in \mathfrak{M}_x(\mathfrak{U})$ be an heir of μ. If ν is not $G(\mathfrak{U})$-invariant there are $\varphi(x; y) \in L(M)$, $a \in \mathfrak{U}$, $g \in G(\mathfrak{U})$ and $r_0, r_1 \in [0, 1]$ such that

$$\nu(\varphi(x; a)) = r_0 > r_1 = \nu(\varphi(g^{-1}x; a))$$

[2] The weak one only differs in requiring the condition for single formulas instead of finitely many.

Letting $r = (r_0 + r_1)/2$ and using the "heir" hypothesis we can contradict $G(M)$-invariance of μ by finding $(b, h) \in M$ that allow to pass from left to right:

$$\begin{array}{ll} \nu(\varphi(x; a)) > r & \mu(\varphi(x; b)) > r \\ \nu(\neg\varphi(g^{-1}x; a)) > 1 - r & \mu(\neg\varphi(h^{-1}x; b)) > 1 - r \\ \nu(g \in G) > 0 & \mu(h \in G) > 0 \end{array}$$ □

Proposition 4.9 ([CPS14, Theorem 2.5])**.** Measures over M have global M-invariant heirs.

Proof. Let $\mu \in \mathfrak{M}_x(M)$ and let $H_\mu \subseteq \mathfrak{M}_x(\mathfrak{U})$ be the set of heirs of μ. Notice that we can write it as the intersection of the closed sets of "Δ-heirs of μ up to ε", where Δ is a finite set of $L(M)$-formulas and $\varepsilon > 0$, namely

$$H_{\mu,\Delta,\varepsilon} = \left\{ \nu \in \mathfrak{M}_x(\mathfrak{U}) \;\middle|\; \forall a \in \mathfrak{U}, \forall (r_\varphi)_{\varphi\in\Delta} \in [0,1)^\Delta \right.$$
$$\left. \Big(\forall b \in M \bigvee_{\varphi\in\Delta} (\mu(\varphi(x; b)) \le r_\varphi - \varepsilon)\Big) \implies \Big(\bigvee_{\varphi\in\Delta} (\nu(\varphi(x; a)) \le r_\varphi)\Big) \right\}$$

Notice that $\{H_{\mu,\Delta,\varepsilon} \mid \Delta \in \mathscr{P}_{\mathrm{fin}}(L(M)), \varepsilon > 0\}$ is closed under finite intersections, so by compactness of $\mathfrak{M}_x(\mathfrak{U})$ we only have to show that each $H_{\mu,\Delta,\varepsilon} \cap \mathfrak{M}^{\mathrm{inv}}_x(\mathfrak{U}, M)$ is non-empty. Fix Δ and ε and invoke Theorem 3.45 to find $p_0, \dots, p_{n-1} \in S(\mu)$ such that for all $b \in M$ and $\varphi \in \Delta$ we have[3]

$$|\mu(\varphi(x; b)) - \mathrm{Av}(p_0, \dots, p_{n-1}; \varphi(x; b))| \le \varepsilon$$

Let $\bar{x} = x_0, \dots, x_{n-1}$, take a completion $\bigcup_{i<n} p_i(x_i) \subseteq p(\bar{x}) \in S_{\bar{x}}(M)$, and then use Corollary 3.93 to find a global M-invariant heir $q(\bar{x})$ of $p(\bar{x})$. Split $q(\bar{x}) = q(x_0, \dots, x_{n-1})$ into n types $q_i(x_i)$ by taking restrictions and then set

$$\nu_{\Delta,\varepsilon} = \mathrm{Av}(q_0, \dots, q_{n-1}) \in \mathfrak{M}_x(\mathfrak{U})$$

We have to check two things:

$\nu_{\Delta,\varepsilon} \in \mathfrak{M}^{\mathrm{inv}}_x(\mathfrak{U}, M)$. By Corollary 3.42 it suffices to show that $\nu_{\Delta,\varepsilon}$ does not fork over M. Suppose that $\psi(x) \in L(\mathfrak{U})$ forks over M and $\nu(\psi(x)) > 0$. Then, by definition, there is $i < n$ such that $\psi(x_i) \in q_i(x_i)$, contradicting the fact that q is M-invariant, hence M-non-forking.

$\nu_{\Delta,\varepsilon} \in H_{\mu,\Delta,\varepsilon}$. Fix $a \in \mathfrak{U}$ and $(r_\varphi)_{\varphi\in\Delta}$ such that

$$\bigwedge_{\varphi\in\Delta} \underbrace{\frac{|i < n \mid \varphi(x_i; a) \in q_i(x_i)|}{n}}_{=\nu_{\Delta,\varepsilon}(\varphi(x;a))} > r_\varphi$$

[3] The fact that Δ is a finite set of formulas instead of a single one can be dealt with by doing definitions by cases, i.e. $\delta(x; y, z) \equiv \bigvee_{\varphi\in\Delta} z = c_\varphi \wedge \varphi(x; y)$.

Since q is an heir of p, by taking a suitable conjunction we can find $b \in M$ such that for all $i < n$ if $\varphi(x_i; a) \in q_i(x_i)$ then $\varphi(x_i; b) \in p_i(x_i)$. Hence

$$\bigwedge_{\varphi \in \Delta} \frac{|i < n \mid \varphi(x_i; b) \in p_i(x_i)|}{n} > r_\varphi$$

By choice of the p_i, for all $\varphi \in \Delta$ this implies $\mu(\varphi(x; b)) > r_\varphi - \varepsilon$. □

Theorem 4.10. All $G(M)$-invariant measures over M can be extended to an M-invariant, $G(\mathfrak{U})$-invariant global measure.

Proof. Take an M-invariant global heir by Proposition 4.9 and apply Proposition 4.8. □

Remark 4.11. It is possible (see [Sim15, Lemma 8.31]) to prove the existence of a $G(\mathfrak{U})$-invariant, definable global extension of a $G(M)$-invariant measure in a more straightforward way[4], basically adapting the proof of Proposition 3.44. The drawback is that this extension will be definable *over some small* $N \succ M$, but we need invariance *over* M in order to produce from it, as we are about to see, a measure finitely satisfiable in M, hence a measure on M^{ext}.

A Retraction

Construction 4.12. Use $N \ {}^{+}\!\!\succ M$ to build L^{ext}, i.e. add to L the predicates $\{R_\varphi(x) \mid \varphi(x) \in L(N)\}$, expand the pair (N, M) to $(N, M, (R_\varphi)_{\varphi \in L(N)})$ and take some $(N', M', (R_\varphi)_{\varphi \in L(N)}) \ {}^{+}\!\!\succ (N, M, (R_\varphi)_{\varphi \in L(N)})$. Let $T^{\text{ext}} = \text{Th}_{L^{\text{ext}}}(M^{\text{ext}})$. Notice that $M^{\text{ext}} \prec^{L^{\text{ext}}} (M', (R_\varphi)_{\varphi \in L(N)})$ and identify $M' \restriction L$ with $\mathfrak{U}$.

Until the end of the section, we will use the notations from the previous construction.

Proposition 4.13 ([CPS14, Proposition 3.9])**.** Let $p(x) \in S^{\text{inv}}_L(M', M)$ and $R_\varphi(x) \in L^{\text{ext}}$. Think of p as a partial L^{ext}-type. If $p(x) \cup R_\varphi(x)$ is consistent, then $p(x) \vdash R_\varphi(x)$.

Proof. If we find $\theta \in p$ such that $\theta(M') \subseteq \varphi(M')$ then, since $(M', (R_\varphi)_{\varphi \in L(N)})$ is a model of $\text{Th}_{L^{\text{ext}}}(M^{\text{ext}})$, we have $\theta \vdash R_\varphi$. To find θ apply Lemma 3.53 translating notations this way:

$$L' := L^{\text{ext}} \cup \{P\} \quad \mathfrak{U} := (N', (R_\varphi)_{\varphi \in L(N)}) \quad p := \{P\}$$
$$\varphi(x; c) := \varphi(x) \quad A := M \quad \theta := \theta$$

Since p is M-invariant and clearly consistent with $P(x)$, and moreover $\theta \in L(M')$, point 3 of the lemma implies that $\theta \in p$, and point 2 says precisely that $\theta(M') \subseteq \varphi(M')$. □

[4] I.e. without passing from the Broom Lemma.

Definition 4.14. In the hypotheses and notations of the previous proposition, let $p' = \{R_\varphi \mid p \vdash R_\varphi\}$.

Remark 4.15. Proposition 4.13 implies that p' is[5] a complete L^{ext}-type.

Definition 4.16. Let $p \in S_x^{\mathrm{inv}}(M', M)$. Define $F_M(p)$ as the "translation" of p' in $L(M')$, i.e.

$$F_M(p) = \{\varphi(x) \in L(M') \mid \exists \psi(x) \in L(N)\ p \vdash R_\psi(x) \text{ and } R_\psi(M) = \varphi(M)\}$$

Remark 4.17. F_M does not depend on N, as long as it codes all externally definable subsets of M.

Lemma 4.18. If $p(x)$ is a global type and $f(x)$ is a definable function, then the pushforward $f^*(p)$ is still a global type. Moreover, if p is M-invariant and f is M-definable, then $f^*(p)$ is M-invariant.

Proof. Since by definition $\varphi(y) \in f^*(p) \iff f^{-1}(\varphi(y)) \in p \iff \varphi(f(x)) \in p$, then $f^*(p)$ is still consistent. Invariance follows from the fact that $\varphi(f(x))$ is a shorthand for $\exists y\ \psi_{y=f(x)}(x, y) \wedge \varphi(y)$, and by hypothesis the previous is an $L(M)$-formula. □

Proposition 4.19 ([CPS14, Proposition 3.10])**.** F_M satisfies what follows.

1. $(F_M(p)) \restriction M = p \restriction M$.
2. It is a continuous retraction $S^{\mathrm{inv}}(\mathfrak{U}, M) \to S^{\mathrm{fs}}(\mathfrak{U}, M)$.
3. It commutes with f^* whenever f is M-definable.

Proof. In the following, suppose that $\psi \in L(N)$ witnesses $\varphi \in F_M(p)$, i.e. $\psi(M) = \varphi(M)$ and $p \vdash R_\psi$.

1. If $\varphi \in L(M)$, then we can set $\varphi = \psi$.
2. We have to check three things.

   It is onto $S^{\mathrm{fs}}(\mathfrak{U}, M)$: Since R_ψ comes from $p' \in S(M^{\mathrm{ext}})$, we have that $\varphi(M) = R_\psi(M) \neq \emptyset$. Since M-finitely satisfiable types are M-invariant, surjectivity is trivial.

   Continuity: By definition $\varphi \in F_M(p)$ if and only if $p \wedge P \vdash \psi$. Hence $F_M^{-1}(\varphi) = \bigcup_{\theta \wedge P \vdash \psi} [\theta]$.

   Fixing $S^{\mathrm{fs}}(\mathfrak{U}, M)$: Suppose that $\neg\varphi \in p$, and let $\theta \in p$ be such that $\theta \vdash R_\psi$. Then p cannot be finitely satisfiable in M, because $\theta(M) \cap \neg\varphi(M) \subseteq \psi(M) \cap \neg\varphi(M) = \varphi(M) \cap \neg\varphi(M) = \emptyset$.

[5] Or, more precisely, "extends uniquely to".

3. We have $\varphi(x) \in F_M(f^*(p)) \Leftrightarrow R_{\psi(x)}(x) \in (f^*(p))'$, and $\varphi(x) \in f^*(F_M(p)) \Leftrightarrow R_\psi(x) \in f^*(p')$, so we might as well work with p'. By compactness and saturation of M', we can find some $B \subseteq M'$ such that $|B| = |N|$, $p \upharpoonright B \vdash p'$ and $f^*(p) \upharpoonright B \vdash (f^*(p))'$. Since B is now small with respect to M', we can find $a \in M'$ such that $a \vDash p \upharpoonright B$. Let $b = f(a)$. On one hand, since $b \vDash f^*(p) \upharpoonright B$, then $b \vDash (f^*(p))'$. On the other hand, since $a \vDash p'$, we have $b \vDash f^*(p')$. □

Example 4.20. Let M be a small model of DLO and let $p = \mathrm{tp}(+\infty/\mathfrak{U})$. Then $F_M(p) = \mathrm{tp}(M^+/\mathfrak{U})$.

Proof. We have to show that $F_M(p) \vdash x < b$ if and only if $b > M$. If $b > M$, then the formula $\varphi(x) = x \geq b$ is such that $\varphi(M) = \emptyset$, so $\neg\varphi \in F_M(p)$. If $b \not> M$, then $\varphi(x) = x > b$ identifies a cut in M that, by saturation, is also identified by $x > c$ for some $c \in N$; setting $\psi(x) = x > c$ proves that $\varphi \in F_M(p)$. □

Extending an Invariant Measure to External Types

By Corollary 3.42 and Proposition 2.104 we can identify respectively $\mathfrak{M}_x^{\mathrm{inv}}(\mathfrak{U}, M)$ and $\mathfrak{M}_x^{\mathrm{fs}}(\mathfrak{U}, M)$ with the space of measures over $S_x^{\mathrm{inv}}(\mathfrak{U}, M)$ and over $S_x^{\mathrm{fs}}(\mathfrak{U}, M)$. This allows to extend F_M from types to measures and generalize Proposition 4.19.

Definition 4.21. $F_M \colon \mathfrak{M}^{\mathrm{inv}}(\mathfrak{U}, M) \to \mathfrak{M}^{\mathrm{fs}}(\mathfrak{U}, M)$ is defined as $F_M(\mu) = F_M^*(\mu)$, where $F_M^*(\mu)$ is the pushforward of μ under $F_M \colon S^{\mathrm{inv}}(\mathfrak{U}, M) \to S^{\mathrm{fs}}(\mathfrak{U}, M)$, i.e. $(F_M^*(\mu))([\varphi(x)]) = \mu(F_M^{-1}([\varphi(x)]))$.

Proposition 4.22 ([CPS14, Proposition 3.15])**.** F_M satisfies what follows.

1. $(F_M(\mu)) \upharpoonright M = \mu \upharpoonright M$.
2. It is a continuous retraction $\mathfrak{M}^{\mathrm{inv}}(\mathfrak{U}, M) \to \mathfrak{M}^{\mathrm{fs}}(\mathfrak{U}, M)$.
3. It commutes with f^* whenever f is M-definable.

Proof. All the verifications boil down to Proposition 4.19. □

Proposition 4.23 ([CPS14, Proposition 3.16])**.** $\mathfrak{M}_x(M^{\mathrm{ext}})$ can be identified with $\mathfrak{M}_x^{\mathrm{fs}}(\mathfrak{U}, M)$.

Proof. The former is the space of measures on $S_x(M^{\mathrm{ext}})$ and the latter is the space of measures on $S_x^{\mathrm{fs}}(\mathfrak{U}, M)$. By Proposition 1.77 together with Corollary 3.58, the two spaces are homeomorphic. □

Remark 4.24. The homeomorphism between $\mathfrak{M}_G(M^{\mathrm{ext}})$ and $\mathfrak{M}_G^{\mathrm{fs}}(\mathfrak{U}, M)$ is also an isomorphism of $G(M)$-flows, i.e. it commutes with multiplication by all the $g \in G(M)$, where $(g \cdot \mu)(\varphi(x)) = \mu(\varphi(g \cdot x))$. Therefore, it preserves $G(M)$-invariance.

Theorem 4.25 ([CPS14, Theorem 3.17 (i)])**.** Every $G(M)$-invariant measure on $S_G(M)$ has an extension to a $G(M)$-invariant measure on $S_G(M^{\text{ext}})$.

Proof. Let $\mu \in \mathfrak{M}_G(M)$ be $G(M)$-invariant, and take a $G(\mathfrak{U})$-invariant $\nu \in \mathfrak{M}^{\text{inv}}(\mathfrak{U}, M)$ by Theorem 4.10. Consider the retraction $F_M(\nu) \in \mathfrak{M}^{\text{fs}}(\mathfrak{U}, M)$. By point 3 of Proposition 4.22 F_M commutes with the pushforward by M-definable functions. Since if $g \in G(M)$ the function $g \cdot -$ is M-definable, and by $G(\mathfrak{U})$-invariance $(g \cdot -^*)(\nu) = \nu$, for all $g \in G(M)$ we have

$$\begin{aligned}(F_M(\nu))(\varphi(g \cdot x)) &= (F_M(\nu))(g^{-1}\varphi(x)) = (F_M(\nu))((g \cdot -)^{-1}\varphi(x)) \\ &= ((g \cdot -)^*(F_M(\nu)))(\varphi(x)) = F_M((g \cdot -)^*(\nu))(\varphi(x)) = (F_M(\nu))(\varphi(x))\end{aligned}$$

To conclude, apply the previous remark. □

There is a number of properties of a definable group that pass to Shelah's expansion, but that we did not even define; the interested reader is addressed to [CPS14]. Also, it was proven in [PYar, Lemma 5.1] that a type is almost periodic if and only if it has an almost periodic external extension.

4.3 Strongly F-Generic Types

In this section we study a characterization of definably amenable groups in NIP[6] theories which was found in [HP11]. The proof presented here will replace [HP11, Remark 5.7] with a result of [CSed] that would be anyway needed in the proof of the Ellis Group Conjecture; as a consequence, Corollary 4.34, follows from Theorem 4.26 instead of being involved in its proof as in [HP11] or [Sim15]. Moreover, since we already proved it, we will use Theorem 4.10 instead of the result mentioned in Remark 4.11.

The main idea is the following. Recall[7] that every compact Hausdorff topological group has a unique left-invariant regular Borel probability measure, called the *Haar measure*, which we will denote with $\mathfrak{h}$. By Theorem 2.59 this is true for G/G^{00}. If G is definably amenable, a left-invariant measure on $S_G(\mathfrak{U})$ can be push-forwarded onto G/G^{00}, and this pushforward must coincide with $\mathfrak{h}$ by uniqueness. Thus, definable amenability of G should be equivalent to the existence of a lifting of $\mathfrak{h}$. It turns out that such a lifting is possible if and only if certain special global types exist. The precise result, around which this section revolves, is the following:

Theorem 4.26. A definable group is definably amenable if and only if it has a global strongly f-generic type.

Let us begin with defining what such a type is.

[6] Recall that Permanent Assumption 4.6 is still valid.

[7] Theorem A.33.

Definition 4.27. A global type $p \in S_G(\mathfrak{U})$ is *strongly f-generic over M* iff for all $g \in G(\mathfrak{U})$ the translate $g \cdot p$ does not fork over M. If we simply say that p is *strongly f-generic* we mean that there is a small M such that p is strongly f-generic over M.

In [HP11] "f-generic" means what here we would call "strongly f-generic over $\emptyset$". Moreover, [Sim15] and a number of papers omit the "strongly". The reason behind is that this diction was first adopted in [CSed], where "f-generic" is reserved for a weaker notion that turned out to allow stronger characterizations. We will not use "plain" f-generic types, but for the benefit of the curious reader let us mention that the difference with Definition 4.27 is that different translates are allowed to be non-forking over different small models, and that even if there are examples of f-generic, non-strongly-f-generic types (see [CSed, Example 3.11]), G has an f-generic if and only if it has a strongly f-generic. See [CSed, Corollary 3.21].

Definition 4.28. If $p \in S_G(\mathfrak{U})$, we let $\operatorname{Stab} p$ denote the stabilizer of p under the action of $G(\mathfrak{U})$, i.e. $\operatorname{Stab} p = \{g \in G(\mathfrak{U}) \mid gp = p\}$.

Proposition 4.29 ([HP11, Proposition 5.6 (i)])**.** The stabilizer of a global strongly f-generic type is G^{00}.

Proof. Let p be strongly f-generic over M.

Claim. $\operatorname{Stab} p = \{g_1^{-1} g_0 \mid g_0 \equiv_M g_1\}$

Proof of the Claim. By Theorem 3.38, for all $g \in G(\mathfrak{U})$ we have $g \cdot p \in S^{\mathrm{inv}}(\mathfrak{U}, M)$. Hence, if $g_0 \equiv_M g_1$, and if $f \in \operatorname{Aut}(\mathfrak{U}/M)$ is such that $f(g_0) = g_1$, we have $g_1 p = f(g_0) f(p) = f(g_0 p) = g_0 p$. This means that $g_1^{-1} g_0 \in \operatorname{Stab} p$. Conversely, let $g \in \operatorname{Stab} p$ and let $h \vDash p \restriction Mg$. Then $\operatorname{tp}(gh/M) = gp \restriction M = p \restriction M = \operatorname{tp}(h/M)$. Set $g_1 = h^{-1} g^{-1}$ and $g_0 = h^{-1}$ and write $g = g_1^{-1} g_0$. □ CLAIM

The Claim proves simultaneously that $\operatorname{Stab} p$ is type-definable over M and that its index is bounded by $|S_G(M)|$. This implies that $\operatorname{Stab} p \supseteq G^{00}$. On the other hand, if $g \notin G^{00}$ then gp is in a different G^{00}-coset than p, and this proves the other inclusion. □

This ensures that the following object is well-defined.

Definition 4.30. If p is strongly f-generic and $\varphi(x) \in L(\mathfrak{U})$, define $U_p \varphi = \{gG^{00} \mid \varphi \in gp\}$.

Proposition 4.31 ([CSed, Proposition 5.1])**.** Every $U_p \varphi$ is constructible.

Proof. Let M be such that p is strongly f-generic over M, let $\pi \colon S_G(\mathfrak{U}) \to G/G^{00}$ be the canonical projection and make the parameters explicit in

$\varphi(x) = \varphi(x; b)$. Write $U_p\varphi = \pi(S)$, where[8] $S = \{g \in G(\mathfrak{U}) \mid \varphi(gx) \in p\}$. Applying Theorem 3.34 to the formula $\psi(x; y, z) = \varphi(z \cdot x; y)$, taking inverse images from $S_{G,y}(M)$ to $S_{G,y}(\mathfrak{U})$ and then setting $y = b$ we can write $S = \bigcup_{n<N}(A_n \wedge \neg B_{n+1})$ for suitable A_n and B_n closed in $S_G(\mathfrak{U})$. Let $A'_n = \pi^{-1}\pi(A_n)$ and $B'_n = \pi^{-1}\pi(B_n)$, which are still closed because π is continuous from a compact space to an Hausdorff one, and let $S' = \bigcup_{n<N}(A'_n \wedge \neg B'_{n+1})$.

Claim. It suffices to prove that $S = S'$.

Proof of the Claim. If $S = S'$, trivially $U_p\varphi = \pi(S) = \pi(S')$. Since, A'_n and B'_n are G^{00}-invariant by definition, i.e. $\pi^{-1}\pi(A'_n) = A'_n$ and $\pi^{-1}\pi(B'_n) = B'_n$, we have that $\pi(A'_n \wedge \neg B'_{n+1}) = \pi(A'_n) \wedge \neg\pi(B'_{n+1})$. But then we can write $U_p\varphi$ as a finite Boolean combination of closed sets in the following way:

$$U_p\varphi = \pi(S') = \pi\left(\bigcup_{n<N}(A'_n \wedge \neg B'_{n+1})\right) = \bigcup_{n<N}(\pi(A'_n) \wedge \neg\pi(B'_{n+1})) \qquad \square_{\text{CLAIM}}$$

($S' \subseteq S$) If $g \in A'_n \wedge \neg B'_{n+1}$ there is $h \in A_n$ such that $h^{-1}g \in G^{00}$. If $h \in B_{n+1}$ then we contradict $g \notin B'_{n+1}$, so we must have $h \in A_n \wedge \neg B_{n+1} \subseteq S$. As p is G^{00}-invariant, S is too by definition, so $g \in S$.

($S \subseteq S'$) Let now $g \in S$ and n be maximal among the ones for which there is an h such that $h^{-1}g \in G^{00}$ and $h \in A_n \wedge \neg B_{n+1}$. By maximality $h \in A'_n \wedge \neg B'_{n+1} \subseteq S'$, and since A'_n and B'_{n+1} are G^{00}-invariant, $g \in S'$. $\square$

Proposition 4.32 ([HP11, Proposition 5.6 (ii)])**.** If G has a global strongly f-generic type, it is definably amenable.

Proof. Fix a global p strongly f-generic over M, and define $\mu(\varphi) = \mathfrak{h}(U_p\varphi)$. This is well-defined because by Proposition 4.29 the stabilizer of p is G^{00} and by Proposition 4.31 U is Borel, hence Haar-measurable. As for $G(\mathfrak{U})$-invariance, since $\mathfrak{h}$ is G/G^{00} invariant,

$$\begin{aligned}\mathfrak{h}(U_p\varphi(gx)) &= \mathfrak{h}(\{hG^{00} \mid \varphi(gx) \in hp\}) = \mathfrak{h}(\{hG^{00} \mid \varphi(x) \in ghp\}) \\ &= \mathfrak{h}(\{ghG^{00} \mid \varphi(x) \in ghp\}) = \mathfrak{h}(U_p\varphi(x)) \quad \square\end{aligned}$$

Thus we have proved half of Theorem 4.26. We now prove a strengthening of the other half.

Proposition 4.33. If G is definably amenable, and M is any small model, then G has a global strongly f-generic type over M.

Proof. If μ is a global $G(\mathfrak{U})$-invariant measure, its restriction to M is $G(M)$-invariant. By Theorem 4.10 there is μ', an extension of $\mu \restriction M$, which is $G(\mathfrak{U})$-invariant and M-invariant; recall that M-invariance is equivalent to

[8] Which differs from $\delta_p\varphi$ as the latter is a subset of $S_G(M)$, while $S \subseteq S_G(\mathfrak{U})$.

M-non-forking by Proposition 3.41. Take any $p \in S(\mu')$. Since μ' is $G(\mathfrak{U})$-invariant, for all $g \in G(\mathfrak{U})$ we have $gp \in S(g\mu') = S(\mu')$, and since μ' is M-invariant, by Proposition 2.104 we have $S(\mu') \subseteq S_x^{\mathrm{inv}}(\mathfrak{U}, M)$. □

Corollary 4.34 ([Sim15, Corollary 8.20])**.** If a group has a global strongly f-generic type then for all small M it has a global type which is strongly f-generic over M.

Proof. By Proposition 4.32 the group is definably amenable. Apply the previous proposition with your favourite small M. □

Remark 4.35. Anyway, it *may* happen that a definably amenable group does not have any global type which is strongly f-generic over $\emptyset$. See [HP11, Proposition 5.11 (ii)].

Notice that strong f-genericity propagates to the orbit closure:

Lemma 4.36. If $p \in S_G(\mathfrak{U})$ is strongly f-generic over M and $q \in \overline{G(\mathfrak{U})p}$, then q is strongly f-generic over M too.

Proof. By continuity[9] of $g \cdot -$, if $q \in \overline{G(\mathfrak{U})p}$ then every gq is still in $\overline{G(\mathfrak{U})p}$. Suppose that $\varphi \in gq$ and φ forks over M. Then since $gq \in \overline{G(\mathfrak{U})p}$ there is h such that $hp \in [\varphi]$, so p cannot be strongly f-generic over M. □

In Definition 4.27 we did not even bother mentioning the word "left", but the reader should understand what we mean if we say "the space S of *right* strongly f-generic global types". Pillay showed in [Pil13, Proposition 2.5] that it is possible to define a product $p \cdot q$ also on S and that this induces and Ellis semigroup structure in which there are no proper ideals and the Ellis groups are isomorphic to G/G^{00}. Anyway, this is *not* a proof of the Ellis Group Conjecture, as S lives inside the $G(\mathfrak{U})$-flow $S_G(\mathfrak{U})$, whereas the Conjecture is about the $G(M)$-flow $S_G(M)$.

4.4 Proof of the Ellis Group Conjecture

We conclude this thesis by studying the proof of the Ellis Group Conjecture from [CSed]. From now on, let G be a group which is $\emptyset$-definable in a NIP theory and suppose it is definably amenable. The work done from Chapter 2 onwards guarantees the following.

Fact 4.37. Up to replacing M with M^{ext} we may assume that all types over M are definable and work directly in $S_G(M)$.

[9] Or, directly, if $gq \in [\varphi]$ then $q \in [g^{-1}\varphi]$, so there is some $hp \in [g^{-1}\varphi]$, hence $ghp \in [\varphi]$.

Proof. Use Corollary 3.58 to trigger Theorem 1.52. By Corollary 3.57, M^{ext} is still NIP, and by Corollary 3.78 and Remark 3.79 G^{00} does not change, nor does the logic topology on G/G^{00}. Definable amenability is preserved by Lemma 4.5 and Theorem 4.25.

Corollary 3.58 ensures that we may work inside $S_G(M^{\text{ext}})$ instead of $S_G^{\text{ext}}(M^{\text{ext}})$, and that the hypotheses of Theorem 1.75 and Theorem 1.79 are still satisfied, so by Corollary 1.76 we have unique coheirs over arbitrary set of parameters and a characterization of the product $p \cdot q$ in terms of the realizations of p and q. □

We still need an important result.

Theorem 4.38 ([CSed, Theorem 5.2 (Baire-Generic Compact Domination)])**.** Let $p \in S_G(\mathfrak{U})$ be strongly f-generic, $\pi\colon S_G(\mathfrak{U}) \to G/G^{00}$ the projection and $\varphi(x) \in L(\mathfrak{U})$. Denote $G = G(\mathfrak{U})$ and define

$$E_p\varphi = \pi([\varphi(x)] \cap \overline{Gp}) \cap \pi([\neg\varphi(x)] \cap \overline{Gp})$$

Then $E_p\varphi$ is closed with empty interior.

Proof. Since π is continuous from a compact space to an Hausdorff one, $E_p\varphi$ is closed. As translating p does not change $\overline{Gp}$, we may assume that $p \in G^{00}$ without changing $E_p\varphi$, and strong f-genericity is preserved by Lemma 4.36. Since by Proposition 4.31 the set $U_p\varphi$ is constructible, and the border of a constructible set has empty interior[10], it suffices to show $E_p\varphi \subseteq \partial U_p\varphi$.

This amounts to proving that, for all $g \in E_p\varphi$ and all open $V \ni g$, both $V \cap U_p\varphi$ and $V \cap (U_p\varphi)^\complement$ are non-empty. Fix such g and V and let us rewrite $U_p\varphi$ as

$$U_p\varphi = \{h \in G/G^{00} \mid hp \in [\varphi]\}$$

Since $g \in E_p\varphi$, there are $q_0 \in [\varphi] \cap \overline{Gp}$ and $q_1 \in [\neg\varphi] \cap \overline{Gp}$ such that $\pi(q_0) = g = \pi(q_1)$. Since $g \in V$ we also have

$$q_0 \in \pi^{-1}V \cap [\varphi(x)] \cap \overline{Gp} \qquad q_1 \in \pi^{-1}V \cap [\neg\varphi(x)] \cap \overline{Gp}$$

Hence, since $\pi^{-1}V \cap [\varphi]$ is a neighbourhood of q_0, it meets Gp, so there is $h_0 \in G(\mathfrak{U})$ such that $h_0p \in \pi^{-1}V \cap [\varphi]$. This implies $\pi(h_0) \in U_p\varphi$ and, trivially, $\pi(h_0p) \in V$. Applying the same argument to q_1 we find $h_1 \in G(\mathfrak{U})$ such that $\pi(h_1) \in U_p\neg\varphi = (U_p\varphi)^\complement$ and $\pi(h_1p) \in V$. To conclude observe that, since we arranged to have $p \in G^{00}$,

$$\pi(h_0) = \pi(h_0p) \in V \cap U_p\varphi \qquad \pi(h_1) = \pi(h_1p) \in V \cap (U_p\varphi)^\complement$$ □

We are almost ready to prove that the Ellis Group Conjecture is true. Injectivity of $\pi \restriction uI$ will be shown through the following lemma.

[10]See Proposition A.5.

Lemma 4.39. If $p_0, p_1 \in uI$ and there is $r \in I$ such that $rp_0 = rp_1$, then $p_0 = p_1$.

Proof. Since $r \in I$, by points 3 and 4 of Theorem 1.32 there is some $v \in \mathrm{Idem}(I)$ such that r belongs to the group vI. Multiplying both sides of $rp_0 = rp_1$ on the left by $r^{-1} \in vI$ and then by u yields $uvp_0 = uvp_1$. By point 2 of Theorem 1.32 $uv = u$, hence $up_0 = up_1$, and since u is the identity of the group uI we have $p_0 = p_1$. □

Theorem 4.40 ([CSed, Theorem 5.6])**.** Let G be a definably amenable NIP group. Then for any minimal ideal I of $S_G(M)$ and idempotent $u \in I$, the natural map $\pi\colon uI \to G/G^{00}$ is injective.

Proof. Given the length of this proof, we break it in three parts. We invite the reader to keep in mind through all of them that, even if we will mention $S_G(\mathfrak{U})$ several times, the Ellis group uI lives inside the $G(M)$-flow $S_G(M)$.

General Set-Up. By Theorem 1.35 we are free to choose the I and $u \in \mathrm{Idem}(I)$ we work with. Using Proposition 4.33 fix $p \in S_G(\mathfrak{U})$ which is strongly f-generic over M. Consider $\overline{G(M)(p \restriction M)}$. It is a subflow of $S_G(M)$, hence it has a minimal subflow I, which is a minimal ideal by Proposition 1.29. By Lemma 4.36 we can assume, up to replacing p with some $\tilde{p} \in \overline{G(\mathfrak{U})p}$ such that $\tilde{p} \restriction M \in I$, that $I = \overline{G(M)(p \restriction M)}$. Fix $u \in \mathrm{Idem}(I)$.

Given $p_0, p_1 \in uI$ such that $\pi(p_0) = \pi(p_1)$ we want to show that $p_0 = p_1$. Since $\mathrm{Ker}\,\pi = G^{00} \cap uI$, we may assume $p_0, p_1 \in G^{00}$. By Lemma 4.39 it suffices to find $r \in I$ such that $rp_0 = rp_1$.

Definitions. Fix any ultrafilter $\mathcal{U} \in \beta(G/G^{00})$ extending the comeagre filter[11]. For all $\varphi(x) \in L(\mathfrak{U})$, the set $E_p\varphi$ is closed with empty interior by Theorem 4.38, hence it is meagre and $E_p\varphi \notin \mathcal{U}$. Define the following objects:

- For each $g \in G/G^{00}$, choose by Theorem 1.83 some $r_g \in I \subseteq S_G(M)$ such that $\pi(r_g) = g$.
- Set $r = \lim_{g \to \mathcal{U}} r_g$, and notice that $r \in I$ because I is closed. We will prove that for this r we have $rp_0 = rp_1$.
- Let $q_0, q_1 \in \overline{G(\mathfrak{U})p} \subseteq S_G(\mathfrak{U})$ extend p_0 and p_1, and notice that they are still in G^{00}.
- Find $a_0 \vDash q_0$ and $a_1 \vDash q_1$ inside a bigger monster $\widetilde{\mathfrak{U}}\,{}^+\!\succ \mathfrak{U}$.
- By Corollary 1.76 every type in $S_G(M)$ has a unique coheir to any bigger set of parameters. Let r' be the unique coheir of r over $\mathfrak{U}a_0a_1$ and, for all $g \in G/G^{00}$, let r'_g be the unique coheir of r_g over $\mathfrak{U}a_0a_1$.
- Find inside $\widetilde{\mathfrak{U}}$ some $b \vDash r'$ and, for all $g \in G/G^{00}$, some $b_g \vDash r'_g$.

[11] See Definition A.6, Remark A.7 and Example A.10.

Computations. Since the subspace of M-finitely satisfiable types is closed, we have $\lim_{g\to\mathcal{U}} r'_g \in S^{\mathrm{fs}}_G(\mathfrak{U}a_0a_1, M)$. Since this type obviously extends $\lim_{g\to\mathcal{U}} r_g = r$, by uniqueness of coheirs we have $r' = \lim_{g\to\mathcal{U}} r'_g$.

We now show that, for $i \in \{0,1\}$,

$$\lim_{g\to\mathcal{U}} \operatorname{tp}(b_g \cdot a_i/\mathfrak{U}) = \operatorname{tp}(b \cdot a_i/\mathfrak{U})$$

Indeed, fix $\varphi(x) \in L(\mathfrak{U})$ and let N be such that $N \succ M$ and $\varphi(x) \in L(N)$. Since r'_g and r' are finitely satisfiable in M, as a special case they are N-invariant. Thus, after finding inside $\mathfrak{U}$ some $c_i \equiv_N a_i$, we can compute[12]

$$\begin{aligned}
\varphi(x) \in \lim_{g\to\mathcal{U}} \operatorname{tp}(b_g \cdot a_i/\mathfrak{U}) &\iff \{g \in G/G^{00} \mid\models \varphi(b_g \cdot a_i)\} \in \mathcal{U}\\
\overset{\dagger}{\iff} \{g \in G/G^{00} \mid\models \varphi(b_g \cdot c_i)\} \in \mathcal{U} &\iff \varphi(x \cdot c_i) \in \lim_{g\to\mathcal{U}} \operatorname{tp}(b_g/\mathfrak{U})\\
\iff \varphi(x\cdot c_i) \in r' \overset{\dagger}{\iff} \varphi(x\cdot a_i) \in r' &\iff \models \varphi(b\cdot a_i) \iff \varphi(x) \in \operatorname{tp}(b\cdot a_i/\mathfrak{U})
\end{aligned}$$

Suppose towards a contradiction that $rp_0 \neq rp_1$, as witnessed by some $\varphi \in L(M)$ such that $\varphi(x) \in rp_0$ and $\neg\varphi(x) \in rp_1$. Then

$$a_i \models q_i \supseteq p_i \qquad b \models r' = r \mid^{\mathrm{ch}}_M \mathfrak{U}a_0a_1 \supseteq r \mid^{\mathrm{ch}}_M Ma_0a_1$$

Hence, by Theorem 1.79, $b\cdot a_i \models rp_i$, and if we set $\theta(x) = \varphi(x\cdot a_0) \wedge \neg\varphi(x\cdot a_1) \in L(Ma_0a_1)$ we have $b \models \theta(x)$. This means that $\{g \in G/G^{00} \mid\models \theta(b_g)\}$ is in $\mathcal{U}$; as a special case, $\{g \in G/G^{00} \mid\models \theta(b_g)\} \cap E_p\varphi^{\complement} \neq \emptyset$, therefore it has a point g. We are now going to show that this is absurd.

Since $g \in (E_p\varphi)^{\complement}$, which is open by Theorem 4.38, and π is continuous, there is $[\psi(x)] \subseteq S_G(\mathfrak{U})$ such that $r'_g \in [\psi]$ and $\pi([\psi]) \cap E_p\varphi = \emptyset$. Since, by choice of g, we have also $\models \theta(b_g)$, then $\psi \wedge \theta \in r'_g$, and by finite satisfiabilty there is $h \in G(M)$ such that $\models \psi(h) \wedge \theta(h)$. Now, $h \models \psi$ implies $\pi(h) \notin E_p\varphi$. On the other hand, since $q_0, q_1 \in G^{00}$, we have $\pi(h \cdot a_0) = \pi(h) = \pi(h \cdot a_1)$. Since $\operatorname{tp}(h \cdot a_i/\mathfrak{U}) = hq_i \in \overline{G(\mathfrak{U})p}$, and $h \models \theta(x)$, i.e. $\models \varphi(h \cdot a_0) \wedge \neg\varphi(h \cdot a_1)$, this shows

$$\pi(h) \in \pi(\underbrace{\overline{G(\mathfrak{U})p} \cap [\varphi(x)]}_{\ni h\cdot q_0}) \cap \pi(\underbrace{\overline{G(\mathfrak{U})p} \cap [\neg\varphi(x)]}_{\ni h\cdot q_1}) = E_p\varphi \qquad \square$$

Corollary 4.41 ([CSed, Corollary 5.7])**.** In a NIP theory, if G is a definably amenable group, then its Ellis group does not depend on the model used to compute it.

[12] † indicates where N-invariance is used.

Appendix A

Appendix

Here we collect some definitions and the statements of some theorems that have been used in the main chapters of this thesis. We usually give a reference to a book, and for a small number of them we also provide proofs. The disclaimer in the Note on References inside the Introduction still applies.

A.1 Combinatorics

Theorem A.1 (Erdös-Rado)**.** Let μ be an infinite cardinal and define inductively $\beth_0(\mu) = \mu$ and $\beth_{n+1}(\mu) = 2^{\beth_n(\mu)}$. Let $[X]^n$ be the set of n-elements subsets of X. Then, for all $n \in \omega$ and all functions $c\colon [(\beth_n(\mu))^+]^{n+1} \to \mu$ there is a subset H of $(\beth_n(\mu))^+$ such that $|H| = \mu^+$ and $c \restriction [H]^{n+1}$ is constant.

Usually, c is called a *colouring of* $[(\beth_n(\mu))^+]^{n+1}$ *with* μ *colours* and H is called an *homogeneous set.* For a proof, see for instance [TZ12, Theorem C.3.2].

A.2 General Topology

Familiarity with basic concepts and definitions is assumed. For instance, we expect the reader to know the definition of product topology and to recognize Proposition A.2 as standard. Any book on basic general topology is a reference, for instance [Bou66].

Proposition A.2. A continuous function from a compact space to an Hausdorff one is closed.

Proposition A.3. Compact Hausdorff spaces are locally compact, i.e. every point has a fundamental system made of compact neighbourhoods.

Definition A.4. A subset of a topological space is *constructible* iff it is a finite Boolean combination of closed sets.

Proposition A.5. In all topological spaces, the border of a constructible set has empty interior.

Proof. By induction on the length of the Boolean combination. If C is closed and $x \in \partial C$, let $V \ni x$ be open and suppose $V \subseteq \partial C$. Since C is closed, $\partial C \subseteq C$. Then V should meet both C and $C^{\complement}$, but since $V \subseteq C$ we have $V \cap C^{\complement} = \emptyset$.

Since $\partial A = \partial(A^{\complement})$, all we are left to do is to prove that if ∂A and ∂B have empty interior then $\partial(A \cup B)$ has empty interior too. Let $x \in \partial(A \cup B)$ and V be a neighbourhood of x. Then V meets $A \cup B$, say it meets A, and also meets $(A \cup B)^{\complement} = A^{\complement} \cap B^{\complement} \subseteq A^{\complement}$. Therefore $\partial(A \cup B) \subseteq \partial A \cup \partial B$. □

Definition A.6. A subset of a topological space is

- *Nowhere dense* iff its closure has empty interior.
- *Meagre* iff it is a countable union of nowhere dense sets.

Remark A.7. Sets which are closed with empty interior are nowhere dense, hence meagre.

A.3 Filters and Ultrafilters

For a thorough treatment of the use of filters in topology, see [Bou66]. Another fine exposition, together with a comparison with *nets*, can be found in [Cla]. For more informations on ultrafilters see for instance [HS98].

Definition A.8. Let $(B, \wedge, \vee, \neg, 0, 1, \sqsubseteq)$ be a Boolean algebra. We call $\mathcal{F} \subseteq B$ a *filter on B* if

- If $x, y \in \mathcal{F}$ then $x \wedge y \in \mathcal{F}$.
- If $x \in \mathcal{F}$ and $x \sqsubseteq y$, then $y \in \mathcal{F}$.
- $0 \notin \mathcal{F}$ and $1 \in \mathcal{F}$.

If we say that $\mathcal{F}$ is a filter on X, but X is not a Boolean algebra, we mean that $\mathcal{F}$ is a filter on the Boolean algebra $\mathscr{P}(X)$.

Remark A.9. Some authors drop the condition $0 \notin \mathcal{F}$ and call the filters that satisfy it *proper*. Notice that if $\mathcal{F}$ satisfies the second condition, but $0 \in \mathcal{F}$, then $\mathcal{F} = B$. For this reason, sometimes B is called the *improper filter*. For us, filters will always be proper unless stated otherwise.

Example A.10. These are some examples of filters.

- For any $0 \neq x \in B$, the family $\mathcal{F} = \{y \in B \mid y \sqsupseteq x\}$ is called the *principal filter on x*.

- The *Fréchet filter* on an infinite set X is the family of cofinite subset of X.
- Let X be a topological space. Then, for any point $p \in X$, the family of neighbourhoods of p is a filter.
- Let X be a topological space. The *comeagre filter* is the filter $\{E \subseteq X \mid E^{\complement} \text{ is meagre}\}$.

Filters are dual to *ideals*, that we are now going to define. The name comes from the fact that, under the identification of Boolean algebras with Boolean rings with unity, the ideals of a Boolean algebra correspond to the proper ideals of the associated Boolean ring.

Definition A.11. Let $(B, \wedge, \vee, \neg, 0, 1, \sqsubseteq)$ be a Boolean algebra. We call $\mathcal{I} \subseteq B$ an *ideal of* B iff

- If $x, y \in \mathcal{I}$ then $x \vee y \in \mathcal{I}$.
- If $x \in \mathcal{I}$ and $x \sqsupseteq y$, then $y \in \mathcal{I}$.
- $1 \notin \mathcal{I}$ and $0 \in \mathcal{I}$.

If we say that $\mathcal{I}$ is an ideal on X, but X is not a Boolean algebra, we mean that $\mathcal{I}$ is an ideal on the Boolean algebra $\mathscr{P}(X)$.

Remark A.12. It is very easy to see that if $\mathcal{F}$ is a filter then $\{\neg x \mid x \in \mathcal{F}\}$ is an ideal, and that if $\mathcal{I}$ is an ideal then $\{\neg x \mid x \in \mathcal{I}\}$ is a filter.

Definition A.13. An *ultrafilter* $\mathcal{U}$ is a filter which is maximal with respect to inclusion, i.e. if $\mathcal{F} \supseteq \mathcal{U}$ is a filter then $\mathcal{F} = \mathcal{U}$.

Proposition A.14. Let $\mathcal{U}$ be a filter. Then the following are equivalent:

1. $\mathcal{U}$ is an ultrafilter.
2. If $x \notin \mathcal{U}$, then $\neg x \in \mathcal{U}$.
3. If $x \vee y \in \mathcal{U}$, then $x \in \mathcal{U}$ or $y \in \mathcal{U}$.

Example A.15. Let $x \in B$ be an *atom* of B, i.e. be such that $x \neq 0$ and if $0 \neq y \sqsubseteq x$ then $y = x$. Then the principal filter on x is an ultrafilter. In an algebra of the form $\mathscr{P}(x)$ the atoms are precisely the singletons $\{x\}$. In this case, the principal ultrafilter on $\{x\}$ is denoted $\sqcup_x$.

Lemma A.16 (Ultrafilter Lemma)**.** Every filter can be extended to an ultrafilter.

Proof. An easy check reveals that the union of a chain of filters is still a filter. Then apply Zorn's Lemma. □

Corollary A.17. If X is an infinite set, there is a non-principal ultrafilter on X.

Proof. Apply the Ultrafilter Lemma to the Fréchet filter. □

Definition A.18. If B is a Boolean algebra, its *Stone space* is the space of ultrafilters on B with the following topology. For all $x \in B$, the set $[x] = \{\mathcal{U} \text{ ultrafilter on } B \mid x \in \mathcal{U}\}$ is a basic open set.

The Stone space of ultrafilters on a set X is denoted βX.

Remark A.19. By Proposition A.14 we have $[x]^\complement = [\neg x]$, so every $[x]$ is clopen and $\{[x] \mid x \in B\}$ is simultaneously a basis for the open sets and a basis for the closed sets.

Remark A.20. It is easy to check that filters on B correspond to closed sets of its Stone space.

Theorem A.21. The Stone space of a Boolean algebra is compact Hausdorff.

Definition A.22. If $f\colon X \to Y$ is a function and $\mathcal{F}$ is a filter on X, the *pushforward* $f^*(\mathcal{F})$ is the ultrafilter on Y defined as

$$U \in f^*(\mathcal{F}) \iff f^{-1}(U) \in \mathcal{F}$$

Remark A.23. The pushforward of an ultrafilter is an ultrafilter.

We recall how to treat convergence in terms of filters. The following definition was already given in Chapter 1.

Definition A.24. If $\mathcal{F}$ is a filter on a topological space Y, and $\ell \in Y$, we say that $\ell = \lim \mathcal{F}$ iff every neighbourhood of ℓ is in $\mathcal{F}$. If $f\colon Z \to Y$ is a function from a set Z to Y and $\mathcal{F}$ is the pushforward $f^*(\mathcal{F}_0)$ of a filter $\mathcal{F}_0$ on Z, we also write $\lim_{z\to\mathcal{F}_0} f(z)$ for $\lim \mathcal{F}$.

Example A.25. The usual limit of a sequence $(a_n)_{n<\omega}$ can be seen as the limit of the pushforward of the Fréchet filter on ω under a_-.

Proposition A.26. Limits on filters commute with continuous functions.

Theorem A.27. A space is compact iff every ultrafilter has at least one limit. A space is Hausdorff iff every ultrafilter has at most one limit.

Proposition A.28. The product topology is the topology of point-wise convergence. In other words, denoting $\pi_j\colon \prod_{i\in I} X_i \to X_j$ the usual projection, an ultrafilter $\mathcal{U}$ on $\prod_{i\in I} X_i$ converges to ℓ if and only if for all $i \in I$, the pushforward $\pi_i^*(\mathcal{U})$ converges to $\pi_i(\ell)$.

Theorem A.29 (Tychonoff)**.** All products of compact spaces are compact.

Let us recall another definition which was already given in Chapter 1.

Definition A.30. If $\mathcal{U} \in \beta X$ and $\mathcal{V} \in \beta Y$, the *tensor product* $\mathcal{U} \otimes \mathcal{V} \in \beta(X \times Y)$ is defined as

$$U \in \mathcal{U} \otimes \mathcal{V} \iff \{x \in X \mid \{y \in Y \mid (x,y) \in U\} \in \mathcal{V}\} \in \mathcal{U}$$

Lemma A.31. Let f be a function with values in a compact Hausdorff space. Then $\lim_{g \to \mathcal{U}} \lim_{h \to \mathcal{V}} f(g,h) = \lim_{(g,h) \to \mathcal{U} \otimes \mathcal{V}} f(g,h) = \lim_{\ell \to f^*(\mathcal{U} \otimes \mathcal{V})} \ell$.

Proof. Let $f_g(h)$ denote $f(g,h)$. On one hand $p = \lim_{(g,h) \to \mathcal{U} \otimes \mathcal{V}} f(g,h)$ if and only if for all open neighbourhoods U of p the following equivalent statements hold:

$$\{g \mid \{h \mid f(g,h) \in U\} \in \mathcal{V}\} \in \mathcal{U} \iff \{g \mid \{h \mid (g,h) \in f^{-1}(U)\} \in \mathcal{V}\} \in \mathcal{U}$$
$$\iff \{g \mid \{h \mid h \in f_g^{-1}(U)\} \in \mathcal{V}\} \in \mathcal{U} \iff \{g \mid U \in f_g^*(\mathcal{V})\} \in \mathcal{U}$$

On the other hand $p = \lim_{g \to \mathcal{U}} \lim_{h \to \mathcal{V}} f(g,h)$ if and only if for all open neighbourhoods U of p we have $\{g \mid \lim_{h \to \mathcal{V}} f_g(h) \in U\} \in \mathcal{U}$. Since by compactness and Hausdorffness $\lim_{h \to \mathcal{V}} f_g(h)$ exists and is unique, by definition and the fact that U is a neighbourhood of each of its points we have $\lim_{h \to \mathcal{V}} f_g(h) \in U \iff U \in f_g^*(\mathcal{V})$, and this proves the first equality. The second one is true by definition. □

Lemma A.32. $\otimes$ is associative.

A.4 Measure Theory and Probability

For the results in this section, see any text covering the foundations of the field, for instance [Hal74].

Theorem A.33 (Haar Measure)**.** Every compact Hausdorff topological group has a unique left-translation-invariant regular Borel probability measure.

Definition A.34. If $f\colon (X_0, \Sigma_0) \to (X_1, \Sigma_1)$ is a measurable function and μ is a measure on (X_0, Σ_0), the *pushforward* $f^*(\mu)$ is the measure on (X_1, Σ_1) defined as $f^*(\mu)(A) = \mu(f^{-1}(A))$.

Proposition A.35. Let $(S, \mathcal{B}, \mu)$ be a probability space. Then there is a probability space (Ω, Σ, P) and a sequence of independent measurable functions $(Y_i\colon \Omega \to S \mid i < \omega)$ all with law μ, i.e. such that $(Y_i)^*(P) = \mu$.

Theorem A.36. Let (Ω, Σ, P) be a probability space, $(X_i\colon \Omega \to \mathbb{R} \mid i < \omega)$ a sequence of independent, equidistributed real random variables in L^1, and $Z_N = \sum_{i<N} X_i$. Then the sequence $((Z_N - \mathbb{E}[Z_N])/N \mid N < \omega)$ converges in probability to 0, i.e. for all $\varepsilon > 0$

$$\lim_{N \to \infty} P\left(\left\{\rho \in \Omega \;\middle|\; \frac{|Z_N(\rho) - \mathbb{E}[Z_N]|}{N} > \varepsilon\right\}\right) = 0$$

A.5 Model Theory

Here we recall some standard theorems of model theory that we will be using from time to time. We omit basic definitions and results; for instance, we will not say what an elementary embedding is or what is the content of the Compactness Theorem. Some reference books are [TZ12, Mar02, Hod93, Hod97, CK90, Poi00].

Definition A.37. Let x be a tuple of variables, $M \vDash T$ and $A \subseteq M$. A *partial type with parameters from* A in variables x is a (non necessarily complete) $L(A \cup x)$-theory. We sometimes say *partial type over* A. A *realization* of a partial type $\pi(x)$ is some $a \in N^{|x|}$, where $N \succ M$, such that $N \vDash \pi(a)$. A *complete type* is a complete partial type. We denote the space of complete types with $S_x(A)$. If we simply say *type*, we mean "complete type".

Remark A.38. A partial (resp. complete) type can be identified with a filter (resp. ultrafilter) on the algebra of $L(A)$-definable sets of $\mathfrak{U}$ in variables x. Therefore, $S_x(A)$ can be identified with the Stone space of $\mathrm{Def}_x(A)$ and this endows it with the topology given in Definition A.18.

Definition A.39. A subset of a structure is *type-definable* iff it can be written as the set of realizations of a partial type. If we say *A-type-definable* or *type definable over A* we mean that such a partial type can be chosen with parameters in A.

Proposition A.40. If $M_0 \prec N$, $M_1 \prec N$ and $M_0 \subseteq M_1$, then $M_0 \prec M_1$.

Proposition A.41 (Amalgamation)**.** If $M_0 \equiv M_1$ there is N such that both M_0 and M_1 can be elementarily embedded[1] in N.

Definition A.42. If κ is an infinite cardinal, a model M is *κ-saturated* iff for all $A \subseteq M$ such that $|A| < \kappa$ and all finite tuples of variables x, every type in $S_x(A)$ is realized in M. We say that M is *strongly κ-homogeneous* iff for all $A \subseteq M$ such that $|A| < \kappa$ all partial elementary maps $A \to M$ extend to automorphisms of M. If every $N \equiv M$ such that $|N| < \kappa$ can be elementarily embedded in M we say that M is *κ-universal*.

Theorem A.43. Every κ-saturated model is κ^+-universal.

Theorem A.44. Let κ be an infinite cardinal. Every model has an elementary extension which is both κ-saturated and strongly κ-homogeneous.

Definition A.45. Let I be an infinite linear order and A a set of parameters. A sequence $(a_i)_{i\in I}$ of tuples from M of the same length is *indiscernible over* A or *A-indiscernible* iff for all $n \in \omega$, all $\varphi(x_0, \dots, x_{n-1}) \in L(A)$ and all $\bar{\imath} = i_0 < \dots < i_{n-1}$ and $\bar{\jmath} = j_0 < \dots < j_{n-1}$ from I we have $M \vDash \varphi(a_{\bar{\imath}}) \leftrightarrow \varphi(a_{\bar{\jmath}})$.

[1] Anyway, it may happen that $M_0 \subseteq M_1$, $M_0 \not\prec M_1$, and $f_0(M_0) \not\subseteq f_1(M_1)$. For instance $(2\mathbb{Z}, <) \subseteq (\mathbb{Z}, <)$, but the embedding is not elementary, and thus by Proposition A.40 in N the inclusion cannot hold.

Definition A.46. Let I be an infinite linear order and $(a_i)_{i\in I}$ be a sequence of tuples from M of the same length ℓ. If A is a set of parameters, the *Ehrenfeucht-Mostowski type of* $(a_i)_{i\in I}$ *over* A is the partial type in the ℓ-tuples of variables $(x_j \mid j < \omega)$ denoted with $\mathrm{EM}((a_i)_{i\in I}/A)$ and defined as follows. For all $\varphi(x_0, \dots, x_{n-1}) \in L(A)$,

$$\varphi(x_0, \dots, x_{n-1}) \in \mathrm{EM}((a_i)_{i\in I}/A) \Leftrightarrow \forall i_0 < \ldots < i_{n-1} \in I \; M \vDash \varphi(a_{i_0}, \ldots, a_{i_{n-1}})$$

Remark A.47. $\mathrm{EM}((a_i)_{i\in I}/A)$ is a complete type if and only if $(a_i)_{i\in I}$ is A-indiscernible.

Lemma A.48 ([TZ12, Lemma 5.1.3 (The Standard Lemma)])**.** Let I and J be two infinite linear orders, $(a_i)_{i\in I}$ a sequence of elements of a structure M and $A \subseteq M$. Then there is $N \succ M$ containing an A-indiscernible sequence $(b_j)_{j\in J}$ such that $\mathrm{EM}((a_i)_{i\in I}/A) \subseteq \mathrm{EM}((b_j)_{j\in J}/A)$.

Theorem A.49 ([Cas11, Proposition 1.6])**.** Let A be a set of parameters, $\kappa > |T| + |A|$ and $\lambda = \beth_{(2^\kappa)^+}$. Let $(b_j \mid j < \lambda)$ be a sequence of tuples all of the same length $\leq \kappa$. Then there is an indiscernible sequence $(a_i \mid i < \omega)$ such that for all $n < \omega$ there are some $j_0 < \ldots < j_{n-1} < \lambda$ with $(a_0, \dots, a_{n-1}) \equiv_A (b_{j_0}, \dots, b_{j_{n-1}})$.

Definition A.50. Let M be an L-structure. We define a language L^{eq} and an L^{eq}-structure M^{eq} as follows. Let $(E_i(x_0, x_1) \mid i \in I)$ be an enumeration of all $\emptyset$-definable equivalence relations on tuples of M, where E_i is a relation on tuples of multi-sort[2] s_i.

Let L^{eq} be L together with a new sort S_i for each $i \in I$ and a function symbol π_i from multi-sort s_i to S_i. Define M^{eq} as M together with, for each $i \in I$, the quotient M^{s_i}/E_i as the interpretation of S_i and the projection $M^{s_i} \to M^{s_i}/E_i$ as the interpretation of π_i.

Proposition A.51 ([TZ12, Proposition 8.4.5])**.** Every $\varphi(\pi_{i_0}(x_0), \pi_{i_1}(x_1))$ is equivalent to some $\psi(x_0, x_1) \in L$.

A.6 The Monster Model

We call $\mathfrak{U}$ a *monster model* for T if it is κ-saturated and κ-strongly homogeneous for a *sufficiently big* κ. What "sufficiently big" means varies among the literature. Some authors require $\mathfrak{U}$ to be proper-class-sized, some assume that κ is inaccessible, the more cautious just regard $\mathfrak{U}$ as a convenient way to simplify statements by avoiding to refer to "some elementary extension" in every theorem and definition.

Since we will sometimes realize types over $\mathfrak{U}$, in order to keep the proofs inside ZFC, for us "sufficiently big κ" means that κ is sufficiently bigger than

[2] I.e., if x_0 has length n, a function from n to the sorts of M.

objects that we call “small”. In other words, we treat $\mathfrak{U}$ as an abuse of notation, and

- if we say “A is small” then $|A| < \kappa$.
- Sometimes, we assume tacitly that also $f(|A|) < \kappa$, where $f(\lambda)$ can be for instance 2^λ, $\beth_n(\lambda)$, or the like.

We now list some conventions and consequences of this. They can be justified via the previous theorems.

- Even if $\mathfrak{U}$ is not uniquely determined, we speak of “the” monster.
- $\vDash \varphi$ means $\mathfrak{U} \vDash \varphi$, and $b \vDash \varphi(x)$ means $\mathfrak{U} \vDash \varphi(b)$.
- All small A are tacitly assumed to be included in $\mathfrak{U}$.
- All small models M are tacitly assumed to be elementary substructures of $\mathfrak{U}$. This implies that if $M \subseteq N$ then automatically $M \prec N$.
- All elementary bijections between small sets can be extended to automorphisms of $\mathfrak{U}$.
- $a \equiv_A b$ if and only if there is $f \in \mathrm{Aut}(\mathfrak{U}/A)$ such that $f(a) = b$.

We would be tempted to add another bullet point stating “all other models are supposed to be small”. Since we sometimes realize global types, or even consider bigger monsters with respect to whom $\mathfrak{U}$ is small, this is simply not true. In order to avoid transforming an abuse of notation into an inconsistency, we do not make this assumption. We would like to ask the reader to forgive us if somewhere in this thesis we have forgotten to say explicitly that an object was supposed to be small, and to try to deduce from the context if this is the case.

Remark A.52. Sometimes definitions and results seem to rely essentially on who $\mathfrak{U}$ or κ are because — say — they quantify on “all small A”, but in fact there are characterizations that avoid mentioning them. See for instance Proposition 2.39.

Acknowledgements

This thesis was built around a beautiful problem, yet allowed me to study several concepts and techniques that are nonetheless interesting in their own right. I do not know if I could have asked for more from a master thesis topic, and I am very grateful to my advisor, Prof. Alessandro Berarducci, for suggesting it, providing a plethora of valuable insights, and because I owe him my interest in model theory and passion for the subject. I would also like to thank Prof. Mauro Di Nasso for reading this thesis now, and for introducing me to ultrafilters some years ago.

I am indebted to a number of my colleagues that either read some part of this thesis, answered my dumb questions, explained to me something, or just sat there listening to me when I needed to explain my problems to someone else. Some of their names, in rigorous alphabetical order, are Alessandro Achille, Agnese Barbensi, Gianluca Basso, Elia Bruè, Jack D'Aurizio, Gianmichele Di Matteo, Sabino Di Trani, Martino Fortuna, Kirill Kuzmin, Luigi Marangio, Filippo Mazzoli, Riccardo Morandin, Giorgio Mossa, Davide Nocera, Roberto Pagaria, Raffaele Scandone, and Andrea Vaccaro. I am sure to be forgetting someone, but obviously I do not remember who.

Writing these pages was the conclusion of a chapter of my life which could not have been possible without the support of my family. Thank you for giving me trust, time, and freedom, and thanks to Nonna Lina for the slaps and the bottles of limoncello.

Listing all the people without whom my last years would have been worse would require several books, and I will not do it. Anyway, if I met you in Aula Studenti, at Radiocicletta, or in Piazza dei Cavalieri and we shared some good laughter, you would probably have been in this hypothetical list, which would also comprehend some people that I was only able to meet very few times in the last years, yet were close to me in one way or another. The list would for sure include most of the names above, plus Wendy Balestri, Matilde Cirini, Alessandro Daluiso, Mauro Daluiso, Francesco Ferrante, Leonardo Fiorini, Marco Fugetto, Vittorio Ghinassi, Simone Lavacca, Daniele Lopez, Greta Malaspina, Davide Nudo, and Gianfranco Tiano.

Last, but not least, I owe a massive amount of gratitude to Federica Ulivieri, but I am unfortunately unable to explain why without reserving a separate chapter for the purpose. Thank you for all you are to me.

Bibliography

[AAG08] E. Akin, J. Auslander and E. Glasner. *The topological dynamics of Ellis actions. Memoirs of the American Mathematical Society*, 913 (2008).

[ABss] A. Achille and A. Berarducci. *Measures* (Work in progress).

[Aus88] J. Auslander. *Minimal Flows and their Extensions*, volume 153 of *Mathematics Studies*. North-Holland (1988).

[Bal88] J. T. Baldwin. *Fundamentals of Stability Theory*, volume 12 of *Perspectives in Mathematical Logic*. Springer-Verlag (1988).

[BBHU08] I. Ben Yaacov, A. Berenstein, C. W. Henson and A. Usvyatsov. *Model theory for metric structures.* In *Model Theory with Applications to Algebra and Analysis*, edited by Z. Chatzidakis, D. Macpherson, A. Pillay and A. Wilkie, volume 2, pages 315–427. Cambridge University Press (2008). Cambridge Books Online.

[BK09] I. Ben Yaacov and H. J. Keisler. *Randomizations of models as metric structures. Confluentes Mathematici*, 1:197–223 (2009).

[Bla93] A. Blass. *Ultrafilters: Where topological dynamics = algebra = combinatorics. Topology Proceedings*, 18:33–56 (1993).

[BN94] A. Borovik and A. Nesin. *Groups of Finite Morley Rank*, volume 26 of *Oxford Logic Guides*. Oxford University Press (1994).

[BOPP05] A. Berarducci, M. Otero, Y. Peterzil and A. Pillay. *A descending chain condition for groups definable in o-minimal structures. Annals of Pure and Applied Logic*, 134:303–313 (2005).

[Bou66] N. Bourbaki. *General Topology*, volume 3 of *Elements of Mathematics*. Addison-Wesley (1966). Translated from the French original.

[BP10] I. Ben Yaacov and A. P. Pedersen. *A proof of completeness for continuous first-order logic. Journal of Symbolic Logic*, 75:168–190 (2010).

[BU10] I. Ben Yaacov and A. Usvyatsov. *Continuous first order logic and local stability. Transactions of the American Mathematical Society*, 362:5213–5259 (2010).

[Cas11] E. Casanovas. *Simple Theories and Hyperimaginaries*, volume 39 of *Lecture Notes in Logic*. Cambridge University Press (2011).

[CK90] C. C. Chang and H. J. Keisler. *Model Theory (3rd edition)*. Elsevier (1990). 3rd edition.

[CK12] A. Chernikov and I. Kaplan. *Forking and dividing in* NTP_2 *theories. Journal of Symbolic Logic*, 77:1–20 (2012).

[Cla] P. L. Clark. *General topology*. http://math.uga.edu/~pete/pointset.pdf.

[CP12] A. Conversano and A. Pillay. *Connected components of definable groups and o-minimality. Advances in Mathematics*, 231:605–623 (2012).

[CPS14] A. Chernikov, A. Pillay and P. Simon. *External definability and groups in NIP theories. Journal of the London Mathematical Society*, 90:213–240 (2014).

[CS13] A. Chernikov and P. Simon. *Externally definable sets and dependent pairs. Israel Journal of Mathematics*, 194:409–425 (2013).

[CSed] A. Chernikov and P. Simon. *Definably amenable NIP groups.* http://arxiv.org/abs/1502.04365 (submitted).

[Ell69] R. Ellis. *Lectures on Topological Dynamics*. Mathematics Lecture Note Series. W.A. Benjamin (1969).

[Gla07a] E. Glasner. *Enveloping semigroups in topological dynamics. Topology and its Applications*, 154:2344–2363 (2007).

[Gla07b] E. Glasner. *The structure of tame minimal dynamical systems. Ergodic Theory and Dynamical Systems*, 27:1819–1837 (2007).

[GPP14] J. Gismatullin, D. Penazzi and A. Pillay. *On compactifications and the topological dynamics of definable groups. Annals of Pure and Applied Logic*, 165:552–562 (2014).

[GPP15] J. GISMATULLIN, D. PENAZZI and A. PILLAY. *Some model theory of* SL$(2, \mathbb{R})$. *Fundamenta Mathematicae*, 229:117–128 (2015).

[Hal74] P. R. HALMOS. *Measure Theory*, volume 18 of *Graduate Texts in Mathematics*. Springer (1974).

[Hod93] W. HODGES. *Model Theory*. Cambridge University Press (1993).

[Hod97] W. HODGES. *A Shorter Model Theory*. Cambridge University Press (1997).

[HP11] E. HRUSHOVSKI and A. PILLAY. *On NIP and invariant measures*. *Journal of the European Mathematical Society*, 13:1005–1061 (2011).

[HPP08] E. HRUSHOVSKI, Y. PETERZIL and A. PILLAY. *Groups, measures and the NIP*. *Journal of the American Mathematical Society*, 21:563–596 (2008).

[HPS13] E. HRUSHOVSKI, A. PILLAY and P. SIMON. *Generically stable and smooth measures in NIP theories*. *Transactions of the American Mathematical Society*, 365:2341–2366 (2013).

[HS98] N. HINDMAN and D. STRAUSS. *Algebra in the Stone-Čech Compactification*, volume 27 of *De Gruyter Expositions in Mathematics*. De Gruyter (1998).

[Las76] D. LASCAR. *Ranks and definability in superstable theories*. *Israel Journal of Mathematics*, 23:53–87 (1976).

[Mar02] D. MARKER. *Model Theory: an Introduction*, volume 217 of *Graduate Texts in Mathematics*. Springer (2002).

[MS94] D. MARKER and C. STEINHORN. *Definable types in o-minimal theories*. *Journal of Symbolic Logic*, 59:185–198 (1994).

[New09] L. NEWELSKI. *Topological dynamics of definable group actions*. *Journal of Symbolic Logic*, 74:50–72 (2009).

[New12] L. NEWELSKI. *Model theoretic aspects of the Ellis semigroup*. *Israel Journal of Mathematics*, 190:477–507 (2012).

[Pie84] J.-P. PIER. *Amenable Locally Compact Groups*. Pure and Applied Mathematics. John Wiley and Sons (1984).

[Pil83] A. PILLAY. *An Introduction to Stabiliy Theory*, volume 8 of *Oxford Logic Guides*. Oxford University Press (1983).

[Pil96] A. PILLAY. *Geometric Stability Theory*, volume 32 of *Oxford Logic Guides*. Oxford University Press (1996).

[Pil13] A. PILLAY. *Topological dynamics and definable groups. Journal of Symbolic Logic*, 78:657–666 (2013).

[Poi83] B. POIZAT. *Paires de structures stables. Journal of Symbolic Logic*, 48:239–249 (1983).

[Poi00] B. POIZAT. *A Course in Model Theory.* Universitext. Springer (2000).

[Poi01] B. POIZAT. *Stable Groups*, volume 87 of *Mathematical Surveys and Monographs*. American Mathematical Society (2001). Translated from the 1987 French original.

[PYar] A. PILLAY and N. YAO. *On minimal flows, definably amenable groups, and o-minimality. Advances in Mathematics* (to appear).

[Sel13] Z. SELA. *Diophantine geometry over groups VIII: Stability. Annals of Mathematics*, 177:787–868 (2013).

[She90] S. SHELAH. *Classification Theory*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North Holland (1990). Revised Edition.

[Sim15] P. SIMON. *A guide to NIP theories*, volume 44 of *Lecture Notes in Logic*. Cambridge University Press (2015).

[TZ12] K. TENT and M. ZIEGLER. *A Course in Model Theory*, volume 40 of *Lecture Notes in Logic*. Cambridge University Press (2012).

[vdD86] L. VAN DEN DRIES. *Tarski's problem and pfaffian functions.* In *Logic Colloquium '84*, edited by J. B. PARIS, A. J. WILKIE and G. M. WILMERS, Studies in Logic and the Foundations of Mathematics, pages 59–90. North Holland (1986).

[vdD98] L. VAN DEN DRIES. *Tame Topology and O-minimal Structures*, volume 248 of *London Mathematical Society Lecture Note Series*. Cambridge University Press (1998).

[Wag00] F. O. WAGNER. *Simple Theories*, volume 503 of *Mathematics and Its Applications*. Kluwer Academic Publishers (2000).

[Wil96] A. J. WILKIE. *Model completeness results for expansions of the ordered field of real numbers by restricted pfaffian functions and the exponential functions. Journal of the American Mathematical Society*, 9:1051–1094 (1996).