Chapter I
Groundwork

The first section of this chapter purports to be a nontechnical introduction to stability theory. It attempts to provide an overview where a neophyte can find his bearings when the subsequent discussions threaten to lose the forest for the trees. The second section contains the necessary list of prerequisites and fixing of notation. We attempt to explain a few of the conventions which have become second nature to specialists but are unfamiliar to many model theorists and unknown to recursion theorists. The third section is an attempt to both summarize the results that a reader should at least be aware of and to show how some of the particular constructions of this book fit into the proof of a more familiar result, Morley's categoricity theorem. One aim of this book is to show the relationship between stability theory and classical algebra. Although various illustrations of this relation will appear in the book, the simplest continuing example is the theory of modules. Thus, in Section 4 of the first chapter we relate the basic properties of the model theory of modules. As we explain in Section 1 the subject of stability theory contains both a structure and a nonstructure side. Most of this book concerns the structure theory. Section 5 contains a brief survey of the methods and results on the nonstructure side.

1. An Overview of Stability Theory

An *algebraic structure theorem* (Ulm's theorem) assigns to each member of a class K of algebras (countable Abelian groups) a system of invariants (the Ulm invariants) which determine the structure up to isomorphism. But many familiar classes, for example, the class of linear orderings, have no such structure theorem. A *classification theorem for universal algebra* determines for each member K of a family $\mathcal{K}$ of classes of algebras whether or not such an algebraic structure theorem holds for K, and, if there is one, exhibits it. This book expounds the fundamentals of what is, perhaps, the first such classification theorem for universal algebra. This theory applies to several families of classes, the most important of which is the family of first order definable classes of algebras or more generally, relational structures.

To aid this exposition and as a step towards the eventual extension of the theory to more general families of classes, we have adopted an axiomatic framework for this book. These axioms describe notions of independence and generation which when satisfied yield the structure theory. Parts A and B of this book describe these axioms and the partition of first order theories which they determine. In Part C we develop the richer part of the structure theory and in Part D use the invariants assigned by the theory to count the number of models of a theory. We outline briefly at the end of Chapter I how to construct the maximum number of nonisomorphic models for each class K which does not have a structure theorem in our sense. However, since a very different and rather set theoretic group of methods is used to prove these results we do not develop them in this book.

The use of the phrase 'system of invariants' in the previous paragraph begs a fundamental question. The usual mathematical meaning of this phrase is a single number (the dimension of a vector space) or a finite sequence of cardinals (the orders of the generators of a finitely generated Abelian group) or an infinite sequence of cardinals (Ulm invariants). We must make a major generalization of this notion which is suggested by the following simple example. Let T be the theory of two equivalence relations E_1, E_2 such that E_1 has infinitely many classes and each E_1 class is refined into infinitely many infinite E_2 classes. A model of T is determined by a set of sets of cardinals: for each E_1 class the set of cardinalities of E_2 classes contained in it. Since it is easy to iterate this example, we are led to the following inductive definition of a system of invariants. A *system of invariants* of rank 0 for a structure A is a set of at most $2^{|T|}$ cardinal numbers each less than or equal to $|A|$. A system of invariants of rank β is a set of at most $2^{|T|}$systems of invariants each of rank less than β. We discuss in the text some natural variants on this idea. The methodological significance of this notion is argued at greater length in [Shelah 1985] and [Baldwin 1986].

There are a number of other uses of the word invariant in mathematics. For example, one refers to the fundamental group of a topological space as an invariant of that space. We demand that a system of invariants for a structure determines that structure up to isomorphism. Still, it would be possible to provide more general systems than suggested here. For example, [Pillay & Steinhorn 1986], one could replace the cardinal numbers by linear orderings. It seems the definition given here produces as simple a set of invariants as one could envision if structures involving refining equivalence relations are to be covered by the theory.

The computation of the number of isomorphism types of members of a class K with cardinality at most $\aleph_\alpha$ provides a challenging test problem for a structure theory. We call the resulting function of α the spectrum function for the class K. We write $I^*(\aleph_\alpha, K) = |\{M \in K : |M| \leq \aleph_\alpha\}|$. Our definition of a system of invariants has the following immediate consequence. If all members of K have a system of invariants of rank less than β, then $I^*(\aleph_\alpha, K) \leq \beth_\beta(\sup(2^{\aleph_0}, |\alpha + \omega|))$. (The $|\alpha + \omega|$ arises here as the number of cardinals less than $\aleph_\alpha$.) A weak version of the contrapositive of this result is even more striking. If, for all sufficiently large κ, K has 2^κ models of power κ then K has no system of invariants of rank β for any β. This raises the question of whether there is a middle ground, a class K which has fewer than the maximal number of models but does not admit a system of invariants. The major conclusion about the spectrum function is that no such middle ground exists. To state this result precisely but simply, let $\mathcal{K}$ be the family of all elementary classes in a countable language. For every $K \in \mathcal{K}$, for all sufficiently large α, $I^*(\aleph_\alpha, K) = 2^{\aleph_\alpha}$ or for all α, $I^*(\aleph_\alpha, K) \leq \beth_\beta(|\alpha + \omega|)$ for some β less than ω_1. This dichotomy is known as 'the main gap'.

The ultimate regularity which arises from the classification of the finite simple groups is the existence of a finite number of families of simple groups. This regularity is obscured by the 26, relatively small, sporadic groups. Similarly, the classification of the spectrum function is obscured by some 'static' among the small models. The correct picture is obtained by first examining the eventual behavior of the spectrum function and then fine-tuning to find its values for small cardinals.

Thus, in this book we will first develop a classification of first order theories. Then for each class of theories we will prove that no structure theory exists for that class or develop one. This classification consists first of a partition of all theories into four stability classes. There are three other, mutually independent, relevant properties: the dimensional order property, the omitting types order property, and depth. A priori, these properties divide all theories into 32 classes. In fact, once a property can be shown to imply that there is no structure theorem we do not have to consider the partition of the theories with that property by the other properties.

In the remainder of this introduction we provide a geography of stable theories. That is, giving few formal definitions and no proofs, we list the types of theories, state the relations between the different types and provide a few examples of each kind. We indicate how invariants are attached to models of theories in the various classes.

In the following survey we deal only with first order theories with *countably* many relation and function symbols. Our ambivalent attitude toward uncountable languages in this book has two sources. First, most of the theories that arise in mathematical practice are naturally formalized in countable languages. However, some are not, notably the theory of vector spaces over the real numbers. Second, a number of important advances in the theory, and I mean important for the study of countable theories, arose from the consideration of uncountable theories. Thus, in the main text, we state most results in a form applicable to uncountable languages and we remark when this case involves added difficulties.

The first dichotomy in our classification of theories is between those which are *stable* and those which are *unstable*. Essentially, a theory T is unstable if it is possible to find in some model of T a linear order which is definable in T on some infinite set of n-tuples. Otherwise, T is stable. Thus, any theory of linear order, any extension of the theory of Boolean algebras, and the theory of arithmetic are all examples of unstable theories. Any unstable theory has 2^λ models of power λ for any uncountable λ ([Shelah 1978]). Thus, no structure theory of the sort described above is possible for the class of models of T.

We will discuss the class of stable theories from an historical viewpoint. The notion of stability arose in the study of categorical theories. A theory T is said to be categorical in power λ if all its models of power λ are isomorphic. Łos [Łos 1954] and Vaught [Vaught 1954] proved in the middle 1950's that if a countable theory with no finite models is categorical in some infinite power λ then it is complete. Łos conjectured in [Łos 1954] that if a countable theory T is categorical in one uncountable power then T is categorical in every uncountable power. (This conjecture is suggested by the Steinitz theorem that every algebraically closed field is determined by the cardinality of its transcendence basis over the prime field.) Morley [Morley 1965] proved the Łos conjecture; in the process, he introduced a wider class of theories than the $\aleph_1$-categorical theories, those that he called totally transcendental and which we now call ω-stable. Examples of ω-stable theories include, in addition to all $\aleph_1$-categorical theories, the theory of differentially closed fields of characteristic zero, the theory of an equivalence relation with infinitely many infinite classes, and the theory of any Abelian group which is the direct sum of a divisible Abelian group and one of bounded exponent. A more specific version of the last example is the theory of the Abelian group $Z_3^\omega \oplus Z_2^\omega$.

While extending Morley's theorem to theories in uncountable languages, Shelah discovered that the crucial properties of ω-stable theories could be generalized by localizing some of Morley's notions from types involving all formulas to ϕ-types, types which contain only instances of a single formula, ϕ. He was then able to isolate a larger class than the ω-stable theories which shared a number of the properties which make ω-stable theories tractable. He called these theories stable. At the same time he discovered a notion which lies between ω-stability and stability which he called the class of superstable theories. In [Shelah 1978], Shelah extends the nonstructure result by proving that any theory which is not superstable has 2^λ models of power λ for every uncountable λ. A natural example of a superstable but not ω-stable theory is the theory of $(Z, +)$. A slightly less natural example of a stable but not superstable theory is the theory of $(Z^\omega, +)$.

The effort to determine exactly which theories have the maximal number of models in each uncountable power required the introduction of two further concepts. A theory is said to have the *dimensional order property* if, speaking roughly and somewhat inaccurately, it is possible to define an ordering within the theory with the use of cardinality quantifiers (There exist κx such that ...). Any such theory has 2^λ models of cardinality λ for any uncountable λ. The simplest example of such a theory is the theory of two crosscutting equivalence relations E_1 and E_2 such that each equivalence class of the relation $E_1 \cap E_2$ is infinite.

Another complication appears exactly at the separation between superstable and ω-stable theories. It is much more difficult to work out the structure theory for superstable theories than it is for ω-stable theories. An interim solution is to restrict the class of models considered when T is only superstable. Of course, such a restriction actually strengthens the result when we prove the restricted class has the maximal number of models. There are a number of families of classes of models which can be treated in a nearly uniform way. The two families $\mathcal{K}$ which receive the most attention in this book are $\mathcal{K}_1 = \{K : K$ is the class of models of an ω-stable theory$\}$ and $\mathcal{K}_2 = \{K : K$ is the class of **S**-models of a superstable theory$\}$. The models in the second family, which we call the strongly ω-saturated models, are also known in the literature as F^a_ω-saturated models, $\aleph_\epsilon$-saturated models and a-models.

The discussion so far shows that if a theory is to have less than the maximal number of models in some uncountable power λ, then T must be superstable and must not have the dimensional order property. In this book we give a complete solution to the spectrum problem for countable ω-stable theories and for S-models of countable superstable theories. The full solution of the original problem, that is the assigning of invariants to the class $\mathcal{K}_3 = \{K :\ K$ is the class of models of a superstable theory$\}$ was obtained recently and will appear in the second edition of [Shelah 1978].

The solution of the spectrum problem for theories in the classes K_1 and K_2 depends upon a key construction which assigns to each model of power λ a skeleton of submodels. Each submodel has cardinality at most 2^ω and the skeleton is partially ordered by the natural tree order on a subset of $\lambda^{<\omega}$. The isomorphism type of the model is determined by the small submodels and this partial ordering. One more property of a theory, which distinguishes between structure and nonstructure, depends on the isomorphism types of trees associated with models of the theory. If one of these trees is not well-founded, the theory is said to be *deep* and has 2^λ models in every uncountable power λ. If not, the theory is *shallow* and the type of structure theory we have described exists. We are able to assign to each such shallow theory a depth β corresponding to the rank of a system of invariants, as discussed above, and to compute the spectrum function of K_1 or K_2 in terms of that depth. Both the dimensional order property and depth are properties which are independent of the stability hierarchy. A good example of a deep ω-stable theory which does not have the dimensional order property is the theory of a single unary function such that each element has infinitely many preimages.

There is a further subdivision of the shallow theories which do not have the dimensional order property. This division is determined by the number of 'bases' to which one can independently assign a dimension. Thus the theory of an equivalence relation with two infinite classes is 2-dimensional since to determine a model we must prescribe the number of elements in each of two classes. Two important classes arise in this context. A theory T is called *bounded* if there is a cardinal $\delta(T)$ such that every model of

T can be determined by specifying less than $\delta(T)$ dimensions. T is called *unbounded* or *multidimensional.* if no such $\delta(T)$ exists. If the theory T is bounded then T has fewer than $|\alpha|^{2^\omega}$ strongly ω-saturated models in power $\aleph_\alpha$. But if T is unbounded, T has at least 2^α strongly ω-saturated models of power $\aleph_\alpha$, where, for simplicity, $\alpha \geq 2^\omega$.

Shelah's solution of the problem for the class K_3 requires the introduction of a number of new ideas which will be treated only peripherally here. The one property necessary to state the final classification is the *omitting types order property.* Again, a non-first order property is used to define a linear ordering. The main gap becomes: If a superstable theory T does not have the dimensional order property, is shallow, and does not have the omitting types order property then $I^*(\aleph_\alpha, T) \leq \beth_\beta(|\alpha|)$ for some β less than ω_1.

The final solution of the spectrum problem can be summarised as follows. Shelah defines an extended logic $L_{\infty,c^+}(Q_D)$ where $c = 2^{\aleph_0}$ and Q_D quantifies over dimensions of independent sets. In [Shelah 198?], Shelah shows that if T is a countable superstable theory without either the dimensional order property or the omitting types order property then each model of T is characterized up to isomorphism by a sentence of $L_{\infty,c^+}(Q_D)$. Since validity in this logic is invariant under changes of the set theoretic universe which do not collapse cardinals and which fix the continuum, the isomorphism classes of such a theory are fairly absolute. If T does not satisfy these conditions, Shelah shows that one can extend the universe of sets by a notion of forcing which preserves cardinals and the continuum yet changes the isomorphism classes of T.

This completes our survey of the relation between the stability hierarchy and the number of models in an uncountable power. We have said nothing as yet about the impact of this theory on the problem of determining the number of countable models of a theory.

Categoricity in power $\aleph_0$ interacts with the stability hierarchy in an interesting way. There are $\aleph_0$-categorical theories which are unstable, for example, dense linear order without endpoints. At the other extreme are the theories which are totally categorical, that is, categorical in every infinite power. Theories in the latter class have a very rich structure which has prompted some of the most sophisticated of the stability theoretic investigations. Every such theory is closely tied to either the theory of an infinite set, the theory of infinite dimensional projective space over a finite field or the theory of infinite dimensional affine space over a finite field. We will summarize this connection and provide more detailed references in Chapter VIII. Lachlan [Lachlan 1974], and independently Shelah [Shelah 1978], proved that a superstable $\aleph_0$-categorical theory is ω-stable. Lachlan's investigation of the stronger conjecture that every $\aleph_0$-categorical stable theory is ω-stable resulted in the definition of a *pseudoplane* and the first introduction of geometry into the subject.

A second major consideration regarding countable models arose from Vaught's conjecture that an $\aleph_1$-categorical theory has either 1 or $\aleph_0$ countable models. Morley [Morley 1970] proved that such a theory has at most $\aleph_0$ countable models. Baldwin and Lachlan [Baldwin & Lachlan 1971] proved this conjecture of Vaught and established that the dimension theory of $\aleph_1$-categorical theories is extremely well behaved. Later Lachlan [Lachlan 1973] extended this result to superstable theories by proving that a countable superstable theory is either $\aleph_0$-categorical or has infinitely many countable models. Lascar [Lascar 1976] and Shelah [Shelah 1978] provided more conceptual proofs of this result. Finally, Pillay [Pillay 1983] gave the simplest and most general version of the theorem. As before, the attempt to extend the result from the superstable case to the stable, that is, to prove that a countable stable theory can not have a finite number of countable models, remains a major open question.

The third major question about countable models, Vaught's conjecture that a first order theory has either $2^{\aleph_0}$ or at most $\aleph_0$ countable models has also been attacked by stability-theoretic means. The only result proved for an arbitrary theory is Morley's theorem that any sentence of $L_{\omega_1,\omega}$ has either $\aleph_1$ or less countable models or exactly $2^{\aleph_0}$ countable models. This proof turns out to be a result in descriptive set theory. This has led to a number of attempts to prove the conjecture itself by the methods of descriptive set theory. So far, there has been no further success on this front. There have been a number of solutions to the problem for special theories such as linear orderings [Rubin 1974] and trees [Steel 1978]. The attempt to solve the problem by breaking it down according to the stability hierarchy has led to the following results. Vaught's conjecture is true for ω-stable theories [Shelah, Harrington, & Makkai 1984]. It is also true for stable theories with Skolem functions [Lascar 1981]. Finally, Vaught's conjecture is true for linearly ordered models with Skolem functions [Shelah 1978a]. The extension of the result in [Shelah, Harrington, & Makkai 1984] to superstable theories seems difficult. Even if it is managed, the usual problem of stable but not superstable theories will remain.

This introduction, and the book as a whole, emphasizes the role of the spectrum problem in the development of stability theory. This makes the organization of the discussion simpler while obscuring some of the other important contributions of the theory. In particular, the structure theory itself may have been lost in this discussion of its application. Many of the structure results in algebra are not actually stated in terms of assigning invariants but rather as the the decomposition of an arbitrary member of some class into a 'product' of known members of the class. The Wedderburn theory is a nice example of this sort. Of course, such a result leads to an assignment of invariants if the 'product' operation is sufficiently tight and the factors are well enough known. We have such a decomposition here, but the 'product' operation is extremely complicated because the factors are partially ordered by a tree. We have chosen to treat this decomposition as an 'internal direct sum'. In [Baldwin & Shelah 1985] it is treated as an 'external direct sum'. That approach emphasizes the significance of the structures $\lambda^{<\omega}$ and λ^{ω}. These are the partially ordered sets of all finite sequences (respectively all countable sequences) from a cardinal λ. The assertion that every sufficiently saturated model with power λ of a stable (superstable) theory can be decomposed over a skeleton of the form λ^{ω} ($\lambda^{<\omega}$) means that the relation between any two elements in M depends on only a countable amount of information even when λ is very large.

There have been major advances in applying these techniques to the study of infinitary languages and languages with generalized quantifiers. Although almost no mention of these matters is made here, the reason for our general approach is the hope of eventually incorporating all these seemingly diverse enterprises into a single framework.

The following table summarises the contents of this section and provides an overview of the stability classification and its relation to the number of countable models. The rows describe a place in the stability classification. A theory belongs in a particular row if it satisfies the condition at the left of the row and does not satisfy the condition in the previous row. The columns indicate the number of countable models. An entry of 'x' means there is no theory of this kind; an entry of '?' means the existence of such a theory is unknown; a numeric entry refers to a theory witnessing that property from the list of theories which follows the table. Many of these examples are discussed in more detail in the text.

An Array of Stable Theories

CLASS	1	$n<\omega$	$\aleph_0$	$\aleph_1$	$\aleph_1<\mu<2^{\aleph_0}$	$2^{\aleph_0}$
$\aleph_1$-categorical	1	x	2	x	x	x
ω-stable	3	x	4	x	x	5
superstable	x	x	6	?	x	7
stable	?	?	8	?	x	9
unstable	10	11	12	?	x	13

1. The theory of an infinite set. The theory of $Z_{p^k}^{\omega}$.
2. The theory of the integers with the successor function. The theory of an algebraically closed field. The theory of $(Q,+)$.
3. The theory of an equivalence relation with two infinite classes.
4. The theory of an equivalence relation with one class of cardinality n for each positive integer n. The theory of an equivalence relation with two classes and a copy of a theory of the second kind in each class.
5. The theory of an equivalence relation with infinitely many classes and a model of a theory of the second kind in each class.
6. The theory of the structure $(2^{\omega}, E_i, +)$ where $+$ denotes coordinatewise addition and $E_i(\sigma,\tau)$ holds if $\sigma|i=\tau|i$.
7. The theory of $(2^{\omega}, E_i)$ with E_i defined as in type 6. The theory of $(Z,+)$.
8. The theory of $(V, B_1, B_2, \ldots)$ where V is a vector space over a finite field and the B_i are a decreasing sequence of subspaces such that

$[B_i : B_{i+1}] = \aleph_0$. (The countable models are determined by the dimension of $\bigcap B_i$.)

9. The theory of (ω^ω, E_i) where E_i is defined as in type 6. The theory of $(Z^\omega, +)$.
10. The theory of dense linear order without endpoints. The theory of atomless Boolean algebras.
11. The theory of the structure $(Q, <, 0, 1, \ldots)$. All known examples are variants of this one.
12. Any disjoint union of an $\aleph_0$-categorical unstable theory with a theory with $\aleph_0$ countable models.
13. Any complete theory of linear orderings which is not $\aleph_0$-categorical [Rubin 1974]. The theory of $(Z, +, <)$.

2. *Basic Notions And Fundamental Conventions*

In this section we fix our notation. We assume the reader is familiar with such results as the compactness and Löwenheim-Skolem theorems and such notions as model completeness and quantifier elimination.

We use α, β, γ for ordinals and κ, μ, λ for cardinals. The notation λ^α generally means the set of functions from α into λ and $\lambda^{<\alpha}$ the set of functions from initial segments of α into λ. This notation should cause little confusion as we rarely use ordinal exponentiation. κ^λ denotes ambiguously the set of functions from λ into κ or the cardinality of that set of functions.

Throughout this book we will be considering models of a complete first order theory T which has no finite models. We denote by $|T|$ the *cardinality* of T, that is, the number of relation symbols, constant symbols, and function symbols in the *language* of T which we denote $L(T)$ or just L. We regard equality as a logical symbol and will always interpret it as identity.

We will write $\phi(\overline{x})$ for the first order formula $\phi(x_1, \ldots, x_n)$ but will not make the n explicit unless its value is crucial. We write M for the structure $(M; \langle R_i^M : R_i \text{ an } L(T) - \text{symbol}\rangle)$. We denote by $\overline{a} \in M$ the assertion that $\overline{a}$ is a finite sequence of elements from M, again suppressing the length of the sequence unless it is particularly relevant. If it is relevant we denote this length by $\lg(\overline{a})$. Technically, such a sequence is a function and we should write $\operatorname{rng}(\overline{a})$; however, we identify such a finite sequence with its range except on those few occasions when confusion is rife. In particular, we frequently write $\overline{a} \cup B$ to denote $\{a_0, \ldots, a_{n-1}\} \cup B$. If α is a function from A to B, we write $\alpha(\overline{a})$ to abbreviate $\langle \alpha(a_1), \ldots, \alpha(a_n)\rangle$. For $\overline{a} \in M$, we write $M \models \phi(\overline{a})$ to indicate that the sequence $\langle a_1, \ldots, a_n\rangle$ satisfies the formula $\phi(\overline{x})$ in the structure M. Let $\phi(\overline{x}; \overline{y})$ be an L-formula and let $\overline{a}$ denote a sequence of elements from M. We denote by $\phi(M; \overline{a})$ the set $\{\overline{m} \in M : \models \phi(\overline{m}; \overline{a})\}$. If $X = \phi(M; \overline{a})$ we say X is *definable over* any set which contains $\overline{a}$. For any structure M, we write $\mathrm{Th}(M)$ for the set of L-sentences true in M. We denote by (M, B), the structure obtained by adding names for each element $b \in B \subsetneq M$. We write $M \models (\exists!^k x)\phi(x)$ if $|\phi(M)| = k$ and $M \models (\exists! x)\phi(x)$ if $|\phi(M)| = 1$.

We will be studying the relation between models $M, N, \ldots$ of the theory T. It is inconvenient if there is an embedding α of M into N such that $M \models \phi(\overline{m})$ but $N \models \neg\phi(\alpha(\overline{m}))$. Of course, this situation does not arise when M is an *elementary submodel* of N, written $M \prec N$. Specifically, this problem cannot arise if the theory T is *model complete.* We also deal with substructures, $A, B, C, \ldots$ of models of T. Again, it is inconvenient if the truth of $\phi(\overline{a})$ depends upon the particular embedding of $\overline{a}$ into a model of T. These problems disappear if we assume that T is *quantifier eliminable.* Fortunately, this is a harmless assumption in our context. For, given any theory T in a language L, we first extend L to a language $L^{\#}$ by adding a relation symbol $R_{\phi(\overline{x})}(\overline{x})$ for each formula $\phi(\overline{x})$. Now we extend T to a theory $T^{\#}$ in $L^{\#}$ by adding the axioms $(\forall \overline{x})[\phi(\overline{x}) \leftrightarrow R_{\phi(\overline{x})}(\overline{x})]$. The resulting theory $T^{\#}$ is equivalent to T for the properties in which we are interested but is quantifier eliminable. For example, for any κ, T and $T^{\#}$ have the same number of models of power κ. Thus we make the following fundamental convention.

2.1 Convention. The theory T is a complete quantifier eliminable theory with no finite models.

While we hold to this convention when proving the general theorems in the text, we will on occasion deal with specific examples of naturally occurring theories that do not admit elimination of quantifiers. Thus it will be essential, although never mentioned, to observe that the definitions we make are preserved in the extension from T to $T^{\#}$.

We may apply any definition made for a theory (e.g. quantifier eliminable) to a structure M by saying M has the property if $\mathrm{Th}(M)$ does.

By means of the translation discussed before Convention 2.1 we are at liberty to assume that all languages are relational. For simplicity of expression we will formulate many examples in languages with function symbols.

For the following exercises we write $M \cap N$ for the intersection of two structures only when the intersection is a common submodel of M and N. This detail will automatically be satisfied on the basis of Convention 2.1. These exercises either are direct translations of terms such as quantifier eliminable or can be deduced from these concepts by the 'method of diagrams'.

2.2 Exercise. If T admits elimination of quantifiers and $\overline{a} \in M \cap N$, where M and N are models of T, show that for any formula $\phi(\overline{x})$, $M \models \phi(\overline{a})$ if and only if $N \models \phi(\overline{a})$.

2.3 Exercise. Show that if T is complete then T has the *joint embedding property.* That is, if M_1 and M_2 are models of T then there is a model N and elementary embeddings α_1 and α_2 of M_1 and M_2, respectively, into N.

2.4 Exercise. Show that if T is model complete then T has the *amalgamation property.* That is, if M_1 and M_2 are models of T and $M_1 \cap M_2 \models T$ then there is a model N and elementary embeddings α_1 and α_2 of M_1 and M_2, respectively, into N which agree on $M_1 \cap M_2$.

2.5 Exercise. Show that if T admits elimination of quantifiers and if M_1, M_2 are models of T with $M_1 \cap M_2 = A$, there is a model N and elementary embeddings α_1 and α_2 of M_1 and M_2, respectively, into N so that α_1 and α_2 agree on A.

Many results in this book are assertions that a certain property holds for all cardinals κ. The proofs given will seem to depend upon the cardinal κ or even upon the existence of a 'well-behaved' model of power κ_1 which is larger than κ. In fact, this dependence is illusory; the arguments are uniform in κ. Rather than continually remark on this uniformity, we assume that there is a well-behaved 'monster' model $\mathcal{M}$ whose cardinality is greater than that of any object in our discussion. For the moment, we take well-behaved to simply mean '$\mathcal{M}$ is a model'. Thus our only use at the moment of this artifice is to be able to write $\mathcal{M} \models \phi$ or $\models \phi$ interchangeably rather than have truth depend on particular models. We have already explained that the assumption that T is quantifier eliminable justifies this simplification. In order to further specify what is meant by well-behaved we recall a few notions from basic model theory.

The Stone representation theorem (e.g. [Sikorski 1964]) associates with each Boolean algebra B a totally disconnected Hausdorff space $S(B)$, the Stone space of B. When B is the Boolean algebra of formulas with parameters from a subset of a model of T, this space has a particularly important interpretation. The analysis of this space is one of the key tools used in this book. We do not prove the Stone representation theorem but we review the principal ingredients of the proof in our context. We begin by defining the Boolean algebra of definable subsets of $\mathcal{M}$ (or $\mathcal{M}^n$). Then we will discuss the natural interpretation of the Stone space of this Boolean algebra. A *formula over* $B \subseteq \mathcal{M}$ is the result of substituting a sequence of parameters $\overline{b}$ from B into an L-formula $\phi(\overline{x};\overline{y})$ to obtain $\phi(\overline{x};\overline{b})$. For any $A \subseteq \mathcal{M}$ we write $\mathrm{Diag}(A)$ for $\{\phi(\overline{a}) : \models \phi(\overline{a})\}$.

2.6 Definition. $F^n(B)$ denotes the Boolean algebra of formulas with n free variables having parameters from B. The Boolean operations in $F^n(B)$ are just the Boolean operations on formulas. Note that we identify formulas up to equivalence in $\mathrm{Th}(\mathcal{M}, B)$.

We write $F(B)$ ambiguously for $F^n(B)$ for any finite n and $F^\omega(B)$ for $\bigcup\{F^n(B) : n < \omega\}$. When we want to emphasize the dependence on T, we write $F(T)$ or $F^n(T)$ for $F(\emptyset)$ or $F^n(\emptyset)$.

The type of $\overline{a}$ over B is a complete description of how $\overline{a}$ relates to the set B. Formally,

2.7 Definition. Let $\lg(\overline{a}) = n$. The (*complete*) *type of* $\overline{a}$ *over* B, denoted $t(\overline{a}; B)$, is

$$\{\phi(\overline{x}) : \models \phi(\overline{a}) \text{ for } \phi \in F^n(B)\}.$$

More generally an α-type p is a consistent set of formulas with α free variables usually denoted $\langle x_i : i < y\alpha\rangle$. We denote by $\operatorname{dom} p$ the *domain of* p, that is, the smallest subset of $\mathcal{M}$ which contains all parameters which occur in a formula in p. We say a type p is *realized* by $\overline{a}$ if $\models \phi(\overline{a})$ for each $\phi(\overline{x}) \in p$. For p a type, we let $p(M)$ equal

$$\{\overline{a} \in M : \models \phi(\overline{a}), \phi \in p\}.$$

If A is a set, by $t(A; B)$ we mean $t(\overline{a}; B)$ where $\overline{a}$ is a fixed enumeration of A. More detail on this convention occurs in Section II.2. The complete type p is a *principal type* if there is a formula $\phi(\overline{x}) \in p$ such that for all $\psi(\overline{x}) \in p$, $\models \phi \rightarrow \psi$.

2.8 Definition. The (nth) *Stone Space of* B, denoted $S^n(B)$, is the collection of all complete n-types over B. We make $S^n(B)$ into a topological space by specifying the collection of sets $U_\phi = \{p \in S^n(B) : \phi \in p\}$ where $\phi \in F^n(B)$ as a basis of open sets.

2.9 Exercise. Show that each set U_ϕ is actually clopen. Conversely, show that each clopen subset of $S^n(B)$ is definable by a single formula with parameters from B.

We write $S(B)$ ambiguously for $S^n(B)$ for any finite n and $S^\omega(B)$ for $\bigcup\{S^n(B) : n < \omega\}$. We sometimes write $S(T)$ for types over the empty set.

An n-*type over* A is a subset of an element of $S^n(A)$. Thus types may be incomplete but are always consistent. If p is a type over A and $B \subseteq A$ then $p|B = \{\phi(\overline{x}; \overline{b}) : \phi(\overline{x}; \overline{b}) \in p \text{ and } \phi(\overline{x}; \overline{b}) \in F(B)\}$. We could regard two types p and q over possibly different domains as equal if they are satisfied by the same elements of the monster model; that is, if $p(\mathcal{M}) = q(\mathcal{M})$. However, this identification destroys the notion of $\operatorname{dom} p$ and makes talk of the cardinality of a type impossible. Thus, we formally view a type as a collection of formulas which has a definite domain and cardinality. Abusing notation we write the expression $t(\overline{a}; A \cup \overline{d}) = t(\overline{a}'; A \cup \overline{d}')$ to mean there is an automorphism of $\mathcal{M}$ which fixes A, maps $\overline{a}$ to $\overline{a}'$ and $\overline{d}$ to $\overline{d}'$.

For $p(\overline{x})$ a type over A, we write $p \vdash \phi(\overline{x})$ if every realization of p satisfies ϕ. Note that this is the same as saying that if L' is formed by adding new constants $\overline{c}$ to L then in L', $\operatorname{Th}(\mathcal{M}, A) \cup p(\overline{c})$ implies $\phi(\overline{c})$. It is tempting to assume that types are closed under logical consequence. Unless this concept is suitably restricted, however, it makes nonsense of the notion of the cardinality of the domain of p. That is, for every p and for every $a \in \mathcal{M}$, $a = a$ would be in p. We sometimes explicitly assume that if $\operatorname{dom} p$ is specified as A, $\phi(\overline{x}; \overline{a})$ is over A and $p \vdash \phi(\overline{x}; \overline{a})$ then $\phi(\overline{x}; \overline{a}) \in p$. Then, types are closed under conjunction.

There is a complete discussion of the impact of various kinds of closure conditions on the definition of type in Chapter II of [Shelah 1978].

2.10 Exercise. Let $A \subseteq B$. Show the restriction map is continuous from $S(B)$ onto $S(A)$.

2.11 Exercise. Show that a point $p \in S(B)$ is isolated just if p is a principal type.

2.12 Exercise. Let (Z, S) denote the integers under successor. Show that $|S^1(Z)| = \aleph_0$ and that it contains a unique non-isolated point. Note that if we replace (Z, S) by $(Z, <)$ then there are two non-isolated points.

2.13 Exercise. Let T be a countable theory. Show that if $|S(A)| = 2^{\aleph_0}$ for some finite A then T has $2^{\aleph_0}$ countable models.

2.14 Exercise. Find an example showing the necessity of the assumption that A is finite in the previous exercise.

Regarding $S(A)$ as a topological space provides a certain insight into the behavior of types. Moreover, it provides a very compact notation for some operations which are model-theoretically important. We next summarize some topological background which is useful in this context.

2.15 Topological Facts. Suppose $A \subseteq B$ and r denotes the restriction map from $S(B)$ onto $S(A)$. Since $S(B)$ is compact, r is continuous, and $S(A)$ is Hausdorff, r is a closed map. Moreover, if X is a closed subset of $S(B)$ then $r(X)$ is compact.

The following definition lists those properties we will need to specify the 'well-behaved' nature of the monster model.

2.16 Definition. i) The model M is *λ-saturated* if for every $A \subseteq M$ with $|A| < \lambda$, every $p \in S(A)$ is realized by some $b \in M$.

ii) The model M is *λ-compact* if for every p with $\operatorname{dom} p \subseteq M$ and $|p| < \lambda$, p is realized in M.

iii) The model M is *λ-homogeneous* if for any subsets A and B of M, such that $|A|, |B| < \lambda$ and $\mathrm{Th}(M, A) = \mathrm{Th}(M, B)$, and for any $m \in M$ there is an $m' \in M$ such that $\mathrm{Th}(M, A, m) = \mathrm{Th}(M, B, m')$.

iv) The structure M is *λ-universal* if for every $N \equiv M$ with $|N| < \lambda$, there is an elementary embedding of N into M.

v) The structure M is *compact, saturated, homogeneous* if it is respectively, $|M|$-compact, $|M|$-saturated, $|M|$-homogeneous.

The following exercises just list some of the basic properties of these notions. They all occur in, for example, [Chang & Keisler 1973].

2.17 Exercise. Show that if M is λ-saturated then M is λ-homogeneous.

2.18 Exercise. Show that if M is homogeneous, $\bar{a}$ and $\bar{b}$ are sequences from M with cardinality less than $|M|$ and $t(\bar{a}; \emptyset) = t(\bar{b}; \emptyset)$ then there is an automorphism of M which takes $\bar{a}$ to $\bar{b}$.

2.19 Exercise. Show that any two saturated elementarily equivalent models with the same cardinality are isomorphic.

2.20 Exercise. Show that any two elementarily equivalent, homogeneous models of the same power which realize the same types (of finite sequences) over the empty set are isomorphic. (This problem is more difficult than the last; it needs to be done by induction on the cardinality of the structures.)

2.21 Exercise. Show that if $\lambda > |T|$ then λ-saturated and λ-compact are equivalent.

2.22 Exercise. Show that a model M is saturated if and only if M is both homogeneous and universal.

2.23 Exercise. Show every saturated model M is $|M|^+$-universal.

2.24 Exercise. Show that any theory T has a κ^+-saturated model of cardinality $\leq 2^\kappa$. (Hint: Choose any model of T with cardinality $\leq 2^\kappa$ and realize every type over each subset of M with cardinality $\leq \kappa$. Repeat this procedure κ^+ times and take the union of the chain of models constructed in this way.)

Assuming the continuum hypothesis one can deduce from Exercise 2.24 that every theory has a saturated model of power $\aleph_1$. Our approach here is to avoid such general set-theoretic hypotheses and search for conditions on a theory which allow a specific model theoretic construction, e.g. a saturated model of power $\aleph_1$. This particular problem is discussed fully in Theorems II.2.20 and III.4.37.

Now we can finish defining the concept of the monster model.

2.25 Convention. We assume that all models and sets that we deal with are contained in a saturated model, called the monster model and denoted $\mathcal{M}$.

As we mentioned earlier, the existence of this model is only assumed as a technical convenience. For those who prefer a more concrete characterization and who believe in large cardinals, we might take $\mathcal{M}$ as a saturated model of power κ where κ is the first strongly inaccessible cardinal.

Items i) and ii) in the following definition are standard notation throughout this book; item iii) fixes our notation for the remainder of this section.

2.26 Notation. i) If α is a partial automorphism of $\mathcal{M}$ and p is a type with $\operatorname{dom} p \subseteq \operatorname{dom} \alpha$ then $\alpha(p)$, sometimes written αp, is the type $\{\phi(\overline{x}; \alpha(\overline{a})) : \phi(\overline{x}; \overline{a}) \in p\}$.

ii) For $A \subseteq M$, $\mathrm{Aut}_A(M)$ denotes the set of automorphisms of M which fix A pointwise. We call an element of $\mathrm{Aut}_A(M)$ an *A-automorphism* of M.

iii) Fix $A \subseteq M \subseteq N$ and suppose M is saturated and $|A| < |M|$. Let r denote the projection map from $S(M)$ onto $S(A)$. Fix $X \subseteq S(M)$; let $P = r(X)$. Let X_N denote $\{\overline{n} \in N : t(\overline{n}; M) \in X\}$ and P_N denote $\{\overline{n} \in N : t(\overline{n}; A) \in P\}$.

Note that $X_N \subseteq P_N$.

For any B, the elements of $S(B)$ can be viewed either syntactically as maximal consistent sets of formulas and thus topologically as elements of the Stone space or semantically as equivalence classes of elements in the monster model. The next series of results show the advantage of adopting the right view at the right time. The following theorem shows that if a subset of $S(M)$ is definable in M but invariant under A-automorphisms of M then, if M is sufficiently saturated, the subset is actually definable over A. Its statement depends on the notation established in 2.26.

2.27 The Basic Definability Theorem. *If X is clopen in $S(M)$ and X is invariant under* $\mathrm{Aut}_A(M)$ *then* $X_N = P_N$.

Proof. Let $\psi(\overline{x};\overline{m})$ define X. By assumption $\psi(\overline{x};\overline{m})$ must be invariant under A-automorphisms of M. Suppose $P_N - X_N$ is not empty; then for some $p \in X$, $q = p|A \cup \{\neg\psi(\overline{x};\overline{m})\}$ is consistent. Let $\overline{a}$ in M realize q and let $\overline{b}$ in M realize $p|A \cup \{\psi(\overline{x};\overline{m})\}$. Then there is an automorphism of M which fixes A and takes $\overline{a}$ to $\overline{b}$. This contradicts the invariance of $\psi(\overline{x};\overline{m})$.

2.28 Corollary. *Under the notation of 2.15 and the conditions of Theorem 2.27:*

i) $r^{-1}r(X) = X$.
ii) $r(X) = P$ *is clopen in* $S(A)$.

Proof. i) simply restates the conclusion of Theorem 2.27. ii) follows by an immediate topological argument noting $r(X)^c = r(X^c)$.

2.29 Exercise. Prove Corollary 2.28 ii) by logical means (a 'double compactness argument').

By taking $M = N = \mathcal{M}$ in Theorem 2.27 we obtain the following corollary which is used extensively in Chapter IV.

2.30 Corollary. *If X is a definable subset of $\mathcal{M}$ and X is fixed by every A-automorphism of $\mathcal{M}$ then X is definable over A.*

2.31 Historical Notes. The basic notions discussed at the beginning of this section can be found in such texts as [Chang & Keisler 1973], [Bridge 1977], [Hodges 198?], [Poizat 1985], [Shoenfield 1967], or [Sacks 1972]. The more advanced material can also be found in [Chang & Keisler 1973] and [Sacks 1972]. Morley pointed out that for the problems discussed here, one could assume without loss of generality that the theory admits elimination of quantifiers. The notion of a monster model was introduced by Shelah.

The vital role of types in model theory was recognized in the middle 1950's. Morley [Morley 1965] introduced the systematic study of Stone spaces. The important role of the basic definability lemma first surfaced in Poizat's proof of the symmetry lemma for forking (cf. exercises III.2.26 and III.2.27.). Its general significance became evident during the 1980-81 model theory year in Jerusalem. Thus, it plays an important role in [Harnik & Harrington 1984]. The formulation here was suggested by Dale Myers.

3. Categoricity Of Countable Theories

In this section we quickly survey the prehistory of stability theory by describing the main results known by 1965 about theories which are categorical in some infinite power. Recall that a theory T is *categorical in power* κ if all models of T which have cardinality κ are isomorphic. To organize this survey we prove Morley's theorem that a countable theory is $\aleph_1$-categorical if and only if it is categorical in every uncountable power. Since this result will appear naturally when we develop the solution to the general spectrum problem, the treatment here is more of an overview. There are several excellent accounts of the theorem in print [Morley 1965], [Chang & Keisler 1973], [Sacks 1972], [Baldwin & Lachlan 1971]; the proof here differs somewhat from all of these. We collect in this section those ingredients of the proof which do not arise in the natural development of the general theory. For other results we will refer to the place in this book where a proof is given. Thus, this survey will provide a point of reference for the more abstract developments later. *We assume in this section that each theory has a countable language.*

Before giving this proof we state a full characterization of countable theories which are $\aleph_0$-categorical. A proof of this characterization can be found in any of the general references cited above.

3.1 Theorem (Ryll-Nardzewski's Theorem). *Let T be a countable complete theory with no finite models. The following are equivalent:*

i) *T is $\aleph_0$-categorical.*
ii) *For each n, $S^n(T)$ is finite.*
iii) *For each n, $F^n(T)$ is finite.*

This proof of Morley's theorem falls into three parts. We first define the concept of a strongly minimal theory. In a strongly minimal theory we are able to directly generalize the notion of algebraic dependence from the theory of vector spaces or fields. Thus, the proof that such a theory is categorical in all uncountable powers just follows the usual argument that the transcendence degree of an algebraically closed field determines the field up to isomorphism. We then show that this result extends to any theory T such that:

i) there is a strongly minimal formula, ϕ, in the language of T.
ii) if M is a model of T, then M is prime and minimal (Definition 3.16) over $\phi(M)$.

It is easily seen that the second condition holds if ϕ is not a two-cardinal formula (see Definition 3.19). Finally, we show that if T is categorical in some uncountable power then it is possible to make a harmless extension of T which satisfies conditions i) and ii). Thus, we can deduce Morley's theorem.

3.2 Definition. i) The formula $\phi(x;\overline{a})$ is *strongly minimal* in T if $\phi(\mathcal{M};\overline{a})$ is infinite and if for every formula $\psi(x;\overline{c})$ (with $\overline{c}$ from the monster model $\mathcal{M}$) either $\phi(\mathcal{M};\overline{a}) \cap \psi(\mathcal{M};\overline{c})$ or $\phi(\mathcal{M};\overline{a}) \cap \neg\psi(\mathcal{M};\overline{c})$ is finite.

ii) The theory T is *strongly minimal* if the formula $x = x$ is strongly minimal in T.

iii) The formula $\phi(x;\overline{a})$ is minimal in M if for any formula $\psi(x;\overline{c})$ with parameters from M either $\phi(M;\overline{a}) \cap \psi(M;\overline{c})$ or $\phi(M;\overline{a}) \cap \neg\psi(Mx;\overline{c})$ is finite.

We call a set A strongly minimal when $A = \phi(M)$ for some strongly minimal formula ϕ.

The first important property of a strongly minimal formula is that it is possible to assign a dimension to the solution set of the formula in any model. We do this by introducing a notion of dependence which is a direct extension of the notion of an element of a field depending on several others.

3.3 Definition. i) The element a is in the *algebraic closure* of the set B if there is a formula $\phi(x;\overline{b})$ with $\overline{b} \in B$ such that $\phi(\mathcal{M},\overline{b})$ is finite and $\models \phi(a,\overline{b})$.

ii) We denote the set of all elements in the algebraic closure of B by $\mathrm{cl}(B)$. If $a \in \mathrm{cl}(B)$, we say $t(a;B)$ is an *algebraic type*; all realizations of such a type are in $\mathrm{cl}(B)$.

iii) A set E is *algebraically independent* if for every $e \in E$, $e \notin \mathrm{cl}(E - e)$.

iv) A set E is a *basis* for a set F if $E \subseteq F$, E is independent, and $\mathrm{cl}(E) \supseteq F$.

We may also describe this situation by saying a depends (algebraically) on B. Note that this dependence notion has finite character. That is, if a depends on B then a depends on a finite subset of B. In Lemma 3.6 we show that within a strongly minimal set this notion of dependence satisfies the axioms for vector space dependence. Chapter II is devoted to a full discussion of the abstract notion of dependence.

3.4 Exercise. Show that if T is $\aleph_0$-categorical and A is finite then $\mathrm{cl}(A)$ is finite.

3.5 Exercise. Show that if $t(a;B)$ is algebraic then $t(a;B)$ is principal. Show the converse fails.

3.6 Lemma. *Suppose $\phi(\mathcal{M})$ is strongly minimal and $(a \cup B \cup C) \subseteq \phi(\mathcal{M})$.*

i) $a \in \mathrm{cl}(a)$.

ii) *If $a \in \mathrm{cl}(B \cup C)$ and $B \subseteq \mathrm{cl}(C)$ then $a \in \mathrm{cl}(C)$.*

iii) *If $a \in \mathrm{cl}(C \cup b)$ and $a \notin \mathrm{cl}(C)$ then $b \in \mathrm{cl}(C \cup a)$.*

iv) *If $a \in \mathrm{cl}(B)$ and $B \subseteq C$ then $a \in \mathrm{cl}(C)$.*

All of these properties except iii) are easily verified by just chasing the definitions. A direct proof of iii) is in [Baldwin & Lachlan 1971]. We will see in Chapter III how to deduce iii) from the general properties of forking. It is now routine to deduce as in linear algebra the invariance of dimension for a strongly minimal set.

3.7 Lemma. If $\phi(M)$ is strongly minimal then all bases X for $\phi(M)$ have the same cardinality.

The second important property of a strongly minimal set is that if $\phi(x;\overline{a})$ is strongly minimal, then for any set A with $\overline{a} \in A$ there is a unique non-algebraic type $p \in S(A)$ with $\phi(x;\overline{a}) \in p$. The next two exercises highlight the ideas needed to prove Theorem 3.10.

3.8 Exercise. If X and Y are independent subsets of a strongly minimal set which have the same cardinality, then $t(X;\emptyset) = t(Y;\emptyset)$. (Hint: Work by induction on $|X|$, showing first that within a strongly minimal set there is a unique non-algebraic type.)

3.9 Exercise. If X and Y realize the same type over the empty set then there is an isomorphism between $\mathrm{cl}(X)$ and $\mathrm{cl}(Y)$. (Hint: Apply Zorn's lemma.)

3.10 Theorem. *If T is a strongly minimal theory then T is categorical in all uncountable powers.*

Proof. Let M and N be models of T and X, Y bases of M and N, respectively. By the first of the above exercises, there is an isomorphism between X and Y; by the second, it extends to an isomorphism between M and N.

We have established the first part of our planned proof. The next step is to extend Theorem 3.10 from strongly minimal theories to those where the universe of each model is determined by the solution set of a strongly minimal formula. For example, the universe might be the algebraic closure of a strongly minimal set. In our outline of the proof we referred to harmless extensions of a theory. The next definition makes precise the class of extensions to which we refer. The exercises show that these extensions are indeed harmless in our context.

3.11 Definition. i) For any theory T, any $M \models T$ and any sequence $\overline{a} \in M$, $T' = \mathrm{Th}(M, \overline{a})$ is an *inessential extension* of T. If $\overline{a}$ is finite, T' is a *finite inessential extension* of T.

ii) If the finite inessential extension T' of T is axiomatized by adding a single sentence to the axioms of T, then T' is a *principal extension* of T.

iii) If some principal extension T' of T has a strongly minimal formula $\phi \in L(T')$ such that if $M \models T'$, $M = \mathrm{cl}(\phi(M))$, the theory T is *almost strongly minimal.*

3.12 Exercise. Show that for any κ with $\kappa^{<\mu} = \kappa$ and any countable theory T, T has a μ-homogeneous model of cardinality κ.

The previous exercise is needed for one implication in the following exercise.

3.13 Exercise. Let T' be a principal extension of the countable theory T. Show that for any uncountable κ, T is κ-categorical if and only if T' is κ-categorical.

3.14 Exercise. Show that if T is almost strongly minimal then T is categorical in every uncountable power.

3.15 Exercise. Show the theory of the Abelian group Z_4^ω is categorical in every uncountable power but not almost strongly minimal. In particular, if $M \models T$ then M is not the closure of the elements of order 2. (Hint: This requires the results on quantifier eliminability of Abelian groups discussed in Section 4.)

We need to extend the notion of a model being generated by a set. The first two notions in the following definition are plausible candidates for the 'model generated by A'. In general, however, if 'generated model' is defined by either of these conditions there is no guarantee that the model generated by A will, in fact, exist. The third notion provides a condition on T which implies that prime models in the sense of Definition 3.16 i) exist. We will see that if T is κ-categorical for some uncountable κ and if A is an infinite definable set, both kinds of model generated by A do exist. Chapter IX contains a detailed study of possible formalizations of the notion of generation.

3.16 Definition. i) The model M is *prime* over $A \subseteq M$ if every elementary map of A into a model N of T extends to an elementary embedding of M into N. If A is empty we simply say M is a prime model of T.

ii) The model M is *minimal* over A if $A \subseteq M$ and there is no proper elementary submodel N of M with $A \subseteq N$.

iii) The theory T is *κ-stable* if for every $A \subseteq \mathcal{M}$ with $|A| \leq \kappa$, $|S(A)| \leq \kappa$.

iv) The theory T is *stable* if it is κ-stable for some κ.

Morley's definition of an ω-stable theory began stability theory. We will see an important source of ω-stable theories in Theorem 3.18 which shows that a countable theory which is categorical in some uncountable power is ω-stable.

The following result demonstrates the powerful effect of combining the existence of strongly minimal formulas and minimal prime models. The lemma generalizes the fact that if $M = \mathrm{cl}(X)$ then M is both prime and minimal over X.

3.17 Lemma. *If some principal extension T' of T has a strongly minimal formula ψ such that each model M of T' is minimal and prime over $\psi(M)$, then T is categorical in all uncountable powers.*

Proof. Let M and N be models of T with cardinality $\kappa > \aleph_0$. Let $\psi(x;\overline{y})$ be an $L(T)$-formula such that $\psi(x;\overline{c})$ is strongly minimal in T' for any $\overline{c}$ which realizes the principal type p. Then we can choose $\overline{c}_1 \in M$ and $\overline{c}_2 \in N$ to expand M and N respectively to models of T'. Then, if X and Y are bases for $\psi(M;\overline{c}_1)$ and $\psi(N;\overline{c}_2)$ respectively, $|X| = |Y| = \kappa$. Since both $\overline{c}_1$ and $\overline{c}_2$ realize p, this implies there is an isomorphism between X and Y which extends, as in the proof of Theorem 3.10, to an isomorphism between

$\psi(M;\overline{c}_1)$ and $\psi(N;\overline{c}_2)$. But since M is minimal and prime over $\psi(M;\overline{c}_1)$, this isomorphism extends further to an isomorphism between M and N.

To complete our proof of Morley's theorem we need to show that if T is categorical in some uncountable power then every model of T is prime and minimal over the interpretation in it of a strongly minimal formula. For 'prime', we quote the following theorem which is proved in Chapter IX. The first part is fairly complicated; the second depends only on the Vaught [Vaught 1961]characterization of prime models.

3.18 Theorem. i) *If T is ω-stable then for every A there is a prime model over A.*
ii) *If M is prime over $\emptyset$ and $\overline{m} \in M$, then* $\mathrm{Th}(M,\overline{m})$ *is a principal extension of T.*

To apply Lemma 3.17 we will show, in fact, that any uncountable model of T, or of any inessential extension of T, is minimal over any infinite definable subset. Moreover, we will find a principal extension of T which has a strongly minimal formula. The next definition and following two theorems accomplish the first aim.

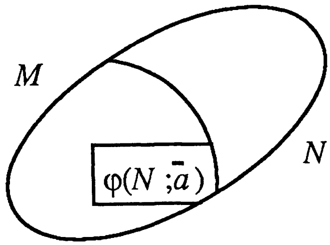

Fig. 1. A two cardinal model

3.19 Definition. (Fig. 1). The formula $\phi(x;\overline{a})$ is a *two-cardinal formula* for T if there exist distinct models $M \subset N$ with $\overline{a} \in M$, such that $\phi(M;\overline{a})$ is infinite and $\phi(M;\overline{a}) = \phi(N;\overline{a})$. Sometimes (M,N) is called a Vaughtian pair.

The next theorem yields that every model of a theory which is categorical in some power is prime and minimal over any definable subset.

3.20 Theorem. *Suppose the countable theory T is categorical in some power $\kappa > \aleph_0$.*

i) *No finite inessential extension of T has a two-cardinal formula.*
ii) *If $M \models T$ and M is prime over $\phi(M;\overline{a})$ then M is minimal over $\phi(M;\overline{a})$.*

Since we show below that a theory which is categorical in some uncountable power is ω-stable and since, in general, ω-stability implies stability Theorem 3.20 i) is easily deduced from the following important result (the two cardinal theorem)proved in Section IX.5.

3.21 Theorem. *If T is a countable stable theory and $\phi(x;\overline{a})$ is a two-cardinal formula in T, then for every $\lambda > \kappa \geq |T|$ there is a model, M, of T with cardinality λ and with $|\phi(M;\overline{a})| = \kappa$.*

Now, to derive Theorem 3.20 i) from Theorem 3.21, suppose that T is κ-categorical. If the finite inessential extension T' of T has a two-cardinal formula, Theorem 3.21 implies there is a model M of T' with cardinality κ but with $|\phi(M;\overline{a})| = \aleph_0$. Clearly there is another model N of T of power κ with $|\phi(N;\overline{b})| = \kappa$ for every choice of the parameters $\overline{b}$. Taking reducts of these two models to the original language contradicts the κ-categoricity of T.

Theorem 3.20 ii) follows immediately from Theorem 3.20 i) and the definition of a two-cardinal formula.

We have thus established condition ii) in our strategy for proving Morley's theorem proclaimed after Definition 3.1. In order to find the strongly minimal formula, required by condition i) of the strategy, we must first show that a theory which is categorical in some uncountable power is ω-stable. This requires the introduction of another powerful technique, the method of Ehrenfeucht-Mostowski models. Surprisingly, this method will play no further role in our development of the positive structure theory. As we indicate in Section 5, it is the principal tool in the construction of many non-isomorphic models. There are actually two ingredients to Ehrenfeucht-Mostowski models, Skolem hulls and indiscernibles. For any first order language L we define the Skolem language associated with L as follows.

3.22 Definition. i) Let $L^0 = L$ and for each $i < \omega$, let L^{i+1} denote the language obtained by adjoining to L^i the $\lg(\overline{y})$-ary function symbol $F_{\exists x \phi}$ for each L^i-formula $(\exists x)\phi(x,\overline{y})$ in which $\overline{y}$ does not occur bound. Let L^{sk}, the *Skolem language* associated with L, denote $\bigcup_{i<\omega} L^i$.

ii) For any L-theory T, the *Skolem theory*, T^{sk}, associated with T is obtained by adding to T the Skolem axioms:

$$(\forall \overline{y})(\exists x)\phi(x,\overline{y}) \rightarrow \phi(F_{\exists x\phi}(\overline{y}),\overline{y})$$

for each L^{sk}-formula $(\exists x)\phi$ in which $\overline{y}$ does not occur bound.

iii) If T is an L-theory and $M \models T$, then a *Skolem expansion*, M^{sk}, of M is any model obtained by interpreting each of the new functions so as to satisfy the Skolem axioms.

3.23 Exercise. For any T, T^{sk} admits elimination of quantifiers.

The next result follows immediately from the previous exercise. Both results are easily proved by induction on the complexity of formulas. See [Chang & Keisler 1973].

3.24 Lemma. *Let M be an L-structure, $X \subseteq M$ and let M_X^{sk} denotes the closure of X under the Skolem functions as interpreted in M^{sk}, then $M_X^{sk} \prec M^{sk}$. We call M_X^{sk} the Skolem Hull of X.*

3.25 Definition. Let $E \subseteq M$ be linearly ordered by $<$, which need not be a definable relation in M. The ordered sequence $(E, <)$ is a *sequence of order indiscernibles* if for any $\phi(\overline{x})$ and any pair $\overline{e}, \overline{e}'$ of finite sequences from E, which are both in increasing order, $\models \phi(\overline{e}) \leftrightarrow \phi(\overline{e}')$.

The following result is proved using the compactness theorem and Ramsey's theorem.

3.26 Theorem (The Ehrenfeucht-Mostowski Theorem). *If T is a theory with infinite models then for every infinite cardinal κ, T has a model which is the Skolem hull of a sequence of order indiscernibles with cardinality κ.*

3.27 Exercise. Prove Theorem 3.26. (Good accounts of this theorem appear in [Chang & Keisler 1973] and [Morley 1967].)

3.28 Exercise. Deduce from (the proof of) Theorem 3.26 that if T has infinite models then for every infinite cardinal κ, T has a model which contains a set of indiscernibles of order type κ.

We deduce from Theorem 3.26, using the last exercise and the properties of Skolem models, the following fundamental result.

3.29 Theorem. *For any theory T and any cardinal κ there is a model M of T with $|M| = \kappa$ such that for any subset A of M with $|A| \leq |T|$, at most $|T|$ types over A are realized in M.*

Proof. Choose by the proof of the last exercise a model, M, of T which is the Skolem hull of a well ordered set, Y, of length κ. Let A be a subset of M with $|A| \leq |T|$. Now each element m of M is determined by a Skolem term $t_m(\overline{x})$ and a finite sequence $\overline{y}_m$ from Y. For any $b \in M$, $t(b; A)$ is determined by the term t_b and the position of $\overline{y}_b$ with respect to the $|A|$ tuples $\overline{y}_a$ for $a \in A$. But there are only $|A|$ ways to interpolate a finite set in a well-ordered set of cardinality κ. Thus, there are at most $(|T| \times |A|) \leq |T|$ types over A realized in M.

Note that Theorem 3.29 does not depend on the countability of T.

3.30 Corollary. *If a countable theory is categorical in some uncountable power then it is ω-stable.*

Proof. Let T be categorical in $\kappa > \aleph_0$. Suppose T is not ω-stable; then there is a countable set A with $|S(A)| > \aleph_0$. By compactness there is a model N of T with $|N| = \kappa$, $A \subset N$ and with uncountably many types over A realized in N. But N must also be isomorphic to the model constructed in Theorem 3.29. This contradiction yields the result.

Many of the basic notions expounded in this book were developed by Shelah in his work on the generalized Łos conjecture: if a theory T is T^+-categorical then it is categorical in all powers. The following exercise shows that such a theory must be $|T|$-stable; in fact, it must be stable in all $\lambda \geq |T|$.

3.31 Exercise. Show that if the theory T (in a possibly uncountable language) is categorical in some $\kappa > |T|$ then T is $|T|$-stable.

We want now to combine 3.30 and 3.20 to show that if T is categorical in some uncountable power we can find a principal extension of T with a strongly minimal formula. We first draw a more precise syntactic consequence of the hypothesis that no finite inessential extension of T has a two-cardinal formula.

3.32 Lemma. *If no finite inessential extension of T has a two-cardinal formula then for every formula $\phi(x;\overline{y})$ there is an integer n such that for every $\overline{c}$, if $|\phi(\mathcal{M};\overline{c})| > n$ then $\phi(\mathcal{M};\overline{c})$ is infinite.*

Proof. Expand L by adding a new predicate symbol U and new constant symbols $\overline{a}$. Let Γ be a set of sentences in the new language which assert:

i) U defines a proper L-elementary submodel of the universe.
ii) Every solution of $\phi(x;\overline{a})$ is in U.
iii) There are infinitely many solutions of $\phi(x;\overline{a})$ in U.

Since no finite inessential extension of T has a two-cardinal model, Γ is inconsistent. But, if for arbitrarily large $m < \omega$, there is a $\overline{c}_m \in \mathcal{M}$ such that $|\phi(\mathcal{M};\overline{c}_m)| = m$, Γ is consistent. Thus, the required n must exist.

To guarantee condition i), denote the relativization of an L-sentence ψ to the predicate U by ψ^U. Place in Γ the sentence $\psi^U \rightarrow \psi$ for each L-sentence ψ.

The following observation is a straightforward consequence of compactness.

3.33 Lemma. *The formula $\phi(x;\overline{a})$ is strongly minimal if and only if for every formula $\psi(x;\overline{y})$ there is a positive integer n such that the two formulas*

$$(\exists^{\geq n}x)[\phi(x;\overline{a}) \wedge \psi(x;\overline{y})]$$

and

$$(\exists^{\geq n}x)[\phi(x;\overline{a}) \wedge \neg\psi(x;\overline{y})]$$

are not both consistent.

One of the basic applications of ω-stability is the construction of minimal sets. The proof is one of a number of variants on the proof that a closed uncountable set of reals contains a perfect subset. A somewhat more detailed explanation of this part of the next argument appears with Theorem III.1.7. We concentrate here on the fact that in the absence of a two cardinal model we can require the set to be strongly minimal.

3.34 Theorem. *If T is an ω-stable countable theory with no two-cardinal formulas then there is a strongly minimal formula with parameters in the prime model of T.*

Proof. Let M be a prime model of T. We can easily choose a formula $\phi(x;\overline{a})$ with parameters from M which is minimal in M (i.e which cannot be split into two infinite parts by a formula with parameters from M. (If not, we would construct an infinite binary tree and contradict ω-stability.) If this formula is not strongly minimal then, by Lemma 3.33 we can find a formula $\psi(x;\overline{y})$ and for each $n \in \omega$ a $\overline{d}_n$ such that both

$$(\exists^{\geq n}x)[\phi(x;\overline{a}) \wedge \psi(x;\overline{d}_n)]$$

and

$$(\exists^{\geq n}x)[\phi(x;\overline{a}) \wedge \neg\psi(x;\overline{d}_n)]$$

hold in $\mathcal{M}$. Now, for arbitrarily large n, either $\phi(\overline{x};\overline{a}) \wedge \psi(\overline{x};\overline{d}_n)$ is infinite and $\phi(\overline{x};\overline{a}) \wedge \neg\psi(\overline{x};\overline{d}_n)$ is finite or vice versa. In the first case $\chi(\overline{x};\overline{y}) = \phi(\overline{x};\overline{y}) \wedge \neg\psi(\overline{x};\overline{y})$ contradicts Lemma 3.32 and the second case is similar.

Now, given a theory T which is categorical in some uncountable power, we can choose by Theorem 3.34 and Theorem 3.20 ii) a principal extension T' which satisfies the hypothesis of Theorem 3.17 and thus is categorical in all uncountable powers. As observed in the Exercises 3.12 and 3.13, this implies T is categorical in all uncountable powers and we have established Morley's theorem.

3.35 Theorem (Morley's Categoricity Theorem). *Let T be a countable first order theory. Then T is categorical in one uncountable power if and only if T is categorical in all uncountable powers.*

The most subtle point of this argument is the choice of the additional parameters for the strongly minimal formula. Let T be an $\aleph_1$-categorical theory. As we noticed in Lemmas 3.6 and 3.7 it is rather easy to see that if $\phi(x)$ is strongly minimal in T then the dependence relation defined by algebraic closure allows us to assign a dimension to the strongly minimal set. If T is not $\aleph_0$-categorical, it is easy to conclude that infinitely many distinct finite dimensions are possible. Thus when ϕ contains no parameters one concludes that T has $\aleph_0$ countable models. Thus, by Theorem 3.34, some finite inessential extension of every $\aleph_1$ but not $\aleph_0$-categorical theory has infinitely many countable models. To prove that T itself has infinitely many models is much more difficult. It may be that different instances of $\phi(x;\overline{y})$ give rise to different dimensions so we cannot assign a dimension to each model. We will establish in Chapter XIII that not only every $\aleph_1$-categorical but every superstable countable theory which is not $\aleph_0$-categorical has $\aleph_0$ countable models.

3.36 Historical Notes. The characterization of $\aleph_0$-categorical theories is due independently to Ryll-Nardzewski [Ryll-Nardzewski 1959], Svenonius [Svenonius 1959], and Engeler [Engeler 1959].

There are at least three fairly distinct proofs in print of Morley's theorem; the one given here is a variant on one of them. The first proof by Morley [Morley 1965] makes substantial use of the rank of types, for example, in the construction of sequences of indiscernibles. The proof in [Chang & Keisler 1973] avoids the use of rank but depends, as does Morley's, on the isomorphism of elementarily equivalent saturated models with the same power. The proof here is closely based on that in [Baldwin & Lachlan 1971] but substitutes the two-cardinal theorem for stable theories (Theorem 3.21) for an *ad hoc* argument in that paper. The one feature which all the proofs have in common is Corollary 3.30, which was first proved by Morley, drawing on the fundamental paper [Ehrenfeucht & Mostowski 1956]. The notion of a strongly minimal set was introduced by Marsh [Marsh 1966]. The same concept was known to Morley as a rank 1 degree 1 set.

The fact that an ω-stable theory with no two-cardinal model has a strongly minimal formula is closely related to the *finite cover property.* This intriguing notion is explored in [Keisler 1967], Chapter II.4 of [Shelah 1978], [Poizat 1983], [Poizat 1984], [Hrushovski 1986], and [Baldwin & Kueker 1980].

4. Introduction to the Model Theory of Modules

One of the major examples used to illustrate the concepts introduced in this book is the theory of modules. In this section we explain the basic facts about the theory of modules which are needed below. Many of the proofs are sketchy; for details see [Ziegler 1984] or [Prest 198?]. By the 'theory of modules' we actually mean the theory of left R-modules for any specified ring R. The following definition makes this precise.

4.1 Definition. The language, L_R, for the theory of R-modules consists of a binary function symbol $+$, a constant symbol 0, and for each element of the fixed ring R a unary function symbol f_r.

The axioms assert that a structure $(M; +, 0, \{f_r : r \in R\})$ is a unitary R-module. That is, $(M; +, 0)$ is an Abelian group, each f_r is an endomorphism of this group, and f_1 is the identity map. We often abbreviate $f_r(y)$ as ry and $(\exists w)rw = y$ by $r|y$. We denote the theory of all R-modules by T_R.

4.2 Examples. The two most important examples are vector spaces over a division ring and modules over the ring of integers Z which are, of course, just Abelian groups. For many of our uses the reader will not suffer by restricting himself to these cases.

4.3 Positive Primitive Formulas. A formula ϕ is a *positive primitive* (p.p.) formula if it is equivalent in T_R to an existential quantification of a conjunction of atomic formulas. In L_R, ϕ is positive primitive just if ϕ has (or is equivalent to a formula with) the form:

$$(\exists \overline{y}) \bigwedge_{i<p} \Big(\sum_{j<m} a_{ij} y_j + \sum_{k<n} b_{ik} x_k = 0 \Big)$$

for some $m, n, p < \omega$. Thus, a positive primitive formula asserts the solvability of a finite system of linear equations. If we write A for the matrix of coefficients (a_{ij}) and B for the matrix of coefficients (b_{ik}) and think of $\overline{y}$ and $\overline{x}$ as column vectors, we can rewrite this p.p. formula as an existentially quantified matrix equation:

$$(\exists \overline{y}) A\overline{y} + B\overline{x} = 0.$$

This allows one to simplify the expression even further when R is a principal ideal domain and, in particular, when $R = Z$ and we are speaking of the integers. We require the following fact found in most graduate algebra texts.

4.4 Theorem. *Suppose H is a square matrix over the principal ideal domain R. Then there are invertible matrices X and Y such that XHY is a diagonal matrix.*

Using this fact it is straightforward to deduce the following normal form for p.p. formulas over a principal ideal domain. Remember that the usual definition of a prime number makes sense in any principal ideal domain.

4.5 Corollary. *If R is a principal ideal domain then every p.p. formula $\phi(\overline{v})$ is equivalent in T_R to a conjunction of formulas of the form $p^k | t(\overline{v})$ or of the form $t(\overline{v}) = 0$ where $t(\overline{v})$ denotes an R-linear combination of variables from $\overline{v}$, p is a prime and k a natural number.*

The following exercises outline a proof from Theorem 4.4 of Corollary 4.5.

4.6 Exercise. Show using Theorem 4.4 that the expression

$$(\exists \overline{y}) A\overline{y} + B\overline{x} = 0$$

can be transformed to

$$(\exists \overline{y}) A'\overline{y} + D\overline{x} = 0 \tag{$*$}$$

where D is a diagonal matrix.

4.7 Exercise. Note that $(*)$ is equivalent to a conjunction of formulas of the form $r | t(\overline{v})$ or of the form $t(\overline{v}) = 0$. Factor the elements r to obtain the form in Corollary 4.5.

Note that the collection of p.p. formulas is closed under existential quantification and conjunction. Moreover, the standard preservation theorems imply that the set of solutions of a p.p. formula is preserved under extension, product and homomorphism of models.

Even more importantly any p.p. formula $\phi(\overline{x})$ with no parameters defines a subgroup of M^n (where $n = \lg(\overline{x})$). A p.p. formula $\phi(\overline{x};\overline{a})$ defines either $\emptyset$ or a coset of the subgroup defined by $\phi(\overline{x};\overline{0})$ where $\overline{0}$ defines a sequence of 0's of appropriate length. To see that $\phi(\overline{x};\overline{0})$ defines a subgroup, note that

$M \models \phi(\bar{0};\bar{0})$ and if $M \models \phi(\bar{a};\bar{0}) \wedge \phi(\bar{b};\bar{0})$ then $M \models \phi(\bar{a}+\bar{b};\bar{0})$. (We write $\bar{a}+\bar{b}$ to denote the coordinatewise sum of $\bar{a}$ and $\bar{b}$.) Summing up, we have:

4.8 Lemma. *If $\phi(\bar{x};\bar{y})$ is a p.p. formula then for any $\bar{a}$ such that $\phi(\bar{x};\bar{a})$ is consistent there is a $\bar{b}$ so that $\phi(M;\bar{a}) = \bar{b} + \phi(M;\bar{0})$.*

4.9 Exercise. Verify Lemma 4.8.

Note that these results imply that any p.p. formula $\phi(\bar{a};\bar{y})$ is either equivalent to a p.p. formula of the form $\phi(\bar{0};\bar{b}-\bar{y})$ or is inconsistent.

4.10 Exercise. Show that if R is commutative and $\phi(\bar{x})$ is a p.p. formula without parameters then $\phi(M)$ is a submodule of M^n.

The p.p. formulas determine the lattice of p.p. definable subgroups where the meet of two such subgroups is their intersection and the join is defined by $(\phi+\psi)(\bar{v}) = (\exists\bar{w})\phi(\bar{w}) \wedge \psi(\bar{v}-\bar{w})$.

The following exercise is crucial to our development.

4.11 Exercise. For any p.p. formula ϕ and any $\bar{a}$, $\bar{b}$, either $\phi(\bar{x};\bar{a})$ and $\phi(\bar{x};\bar{b})$ are equivalent or they are contradictory.

Here is the main theorem of this section.

4.12 Theorem. *Fix an R-module M. For every formula $\phi(\bar{x})$ there is a formula $\phi^*(\bar{x})$ which is a Boolean combination of p.p. formulas such that $M \models \phi(\bar{x}) \leftrightarrow \phi^*(\bar{x})$.*

The proof of this result depends on one algebraic result and one easy combinatorial lemma.

4.13 Lemma ([Neumann 1954]). *Suppose H_i for $i < l$ with $l < \omega$ are Abelian groups. Fix $k \leq l$ such that for each $i < k$, $[H_0 : H_0 \cap H_i]$ is finite while for $k \leq i < l$, $[H_0 : H_0 \cap H_i]$ is infinite. If $H_0 + a_0 \subseteq \bigcup_{i<l} H_i + a_i$ for some sequence $\langle a_i : i < l\rangle$, then $H_0 + a_0 \subseteq \bigcup_{i<k} H_i + a_i$.*

4.14 Exercise (A combinatorial lemma). Let A and $A_0, \ldots, A_{n-1}$ be finite sets. Show $A \subseteq \bigcup_{i<n} A_i$ iff

$$\sum_{\Delta \subseteq n} (-1)^{|\Delta|} |A \cap \bigcap_{i \in \Delta} A_i| = 0.$$

Proof of Theorem 4.12. We will proceed by induction on quantifiers. Thus, consider the formula $(\forall y)\psi(\bar{x};y)$ and suppose by induction that ψ is a Boolean combination of p.p. formulas. Since universal quantification distributes over conjunction and p.p. formulas are closed under conjunction we may assume ψ has the form:

$$\phi_0(\bar{x};y) \rightarrow \bigvee_{1 \leq i < n} \phi_i(\bar{x};y)$$

where each ϕ_i is p.p. For each $i < n$, let H_i denote the subgroup $\phi_i(\bar{0};M)$. Then for any $\bar{b}$, $\psi(\bar{b};y)$ is equivalent (by the remark after Lemma 4.13) to

$$H_0 + a_0 \subseteq \bigcup_{i<n} H_i + a_i$$

for appropriate choice of the a_i. Note that some of the $H_i + a_i$ may be empty. If we reorder the a_i so that $[H_0 : H_0 \cap H_i] < \omega$ iff $i < k$ then apply Lemma 4.8, we see that $\psi(\overline{a}; y)$ is equivalent to:

$$H_0 + a_0 \subseteq \bigcup_{i<k} H_i + a_i. \tag{$*$}$$

Now the key fact is that the truth of $(*)$ can be expressed by a formula which does not mention the a_i. To see this let H denote $\bigcap\{H_i : i < k\}$ and let for $\Delta \subseteq k - \{0\}$

$$N_\Delta = [H_0 \cap \bigcap_{i\in\Delta} H_i : H].$$

Let $m = [H_0 : H]$. Then $H_0 + a_0$ is the union of m cosets of H. For each i, if $(H_i + a_i) \cap (H_0 + a_0) = B_i$ is not empty then B_i contains $n_i = [H_0 \cap H_i : H]$ cosets of H. Thus $H_0 + a_0 \subseteq \bigcup_{i<k} H_i + a_i$ if the number of distinct cosets of H which occur in $\bigcup\{B_i : i < k\}$ is greater than m. The number of these cosets is not $\Sigma\ n_i$ (since this counts various cosets more than once) but must be computed using the N_Δ and the combinatorial lemma. That is, if A_i is the collection of cosets of H in B_i, then $(*)$ is equivalent to $A_0 \subseteq A_1 \cup A_2 \cup \ldots \cup A_{k-1}$. By the combinatorial lemma, $(*)$ is equivalent to

$$\sum_{\Delta\subseteq\Gamma(\overline{a})} (-1)^{|\Delta|} N_\Delta = 0$$

where $\Gamma(\overline{a}) = \{i \in k - \{0\} : B_i{\neq}\emptyset\}$.

Now to remove the apparent dependence of $\Gamma(\overline{a})$ on $\overline{a}$, note that $\Gamma(\overline{a})$ is a subset of $k - \{0\}$ and the statement

$$\sum_{\Delta\subseteq\Gamma(\overline{a})} (-1)^{|\Delta|} N_\Delta = 0$$

depends not on $\overline{a}$ but only on this subset. Let Λ be the collection of subsets Γ of $2^{k-\{0\}}$ such that: $\sum_{\Delta\subseteq\Gamma}(-1)^{|\Delta|}N_\Delta = 0$. Let

$$\phi_\Gamma = (\exists y)\phi_0(\overline{x}; y) \wedge \bigwedge_{\substack{i\in\Delta\\ \Delta\in\Gamma}} \phi_i(\overline{x}; y).$$

Then $(\forall y)\psi(\overline{x};\mathrm{y})$ is equivalent to

$$\bigwedge_{\Gamma\in\Lambda} \phi_\Gamma \ \wedge \ \bigwedge_{\Gamma\notin\Lambda} \neg\phi_\Gamma$$

which is a Boolean combination of p.p. formulas.

Note that the definition of the Boolean combination of p.p. formulas which is equivalent to ψ depends only on the N_Δ and these numbers are the same for elementarily equivalent models. Conversely, if for each Δ, $N_\Delta(M) = N_\Delta(N)$ then M and N are elementarily equivalent. With some further notation we rephrase this result in a more elegant form.

4.15 Notation. If ϕ and ψ are p.p. formulas without parameters and N is a module let $n(\phi/\psi; N) = [\phi(N) : \phi(N) \cap \psi(N)]$. Note that $n(\phi/\psi; N)$ is a natural number or ∞.

For every ϕ, ψ and n, there is a sentence $\chi_{\phi,\psi,n}$ such that for any module N, $N \models \chi_{\phi,\psi,n}$ if and only if $n(\phi/\psi; N) \geq n$. Following [Prest 198?], we call any sentence which expresses these invariants of a module an *invariants sentence.*

4.16 Corollary. *Suppose N and M are R-modules and for each pair of p.p. formulas without parameters, ϕ and ψ, $n(\phi/\psi; M) = n(\phi/\psi; N)$, then $M \equiv N$.*

4.17 Exercise. Write out the general form of an invariants sentence. Notice that it is in AE form. Conclude that if M and N are modules and $M \equiv_{AE} N$ then $M \equiv N$.

A standard compactness argument solves the following exercise ([Prest 198?]).

4.18 Exercise. *Show that for any ring R and any formula $\chi(\overline{x})$ there is a p.p. formula $\theta(\overline{x})$ and an invariants sentence σ such that*

$$T_R \vdash (\forall \overline{x})[\chi(\overline{x}) \leftrightarrow \sigma \wedge \theta(\overline{x})].$$

4.19 Pure Submodules. The submodule M of N is *pure* in N if every p.p. formula with parameters from M which is satisfiable in N is satisfiable in M. Of course, this is a weakening of the notion of elementary submodel but in the light of Theorem 4.13 it is not a very great weakening. In fact, it is easy to see:

4.20 Lemma. *If $M \subseteq N$ and $M \equiv N$ then $M \prec N$ iff M is pure in N.*

4.21 Exercise. Show that if M is a direct factor of N then M is pure in N.

Derive the next exercise from the last one.

4.22 Exercise. Show that every subspace of a vector space is a pure submodule.

4.23 Purity in Abelian Groups. For Abelian groups or, more generally, modules over principal ideal domains, the notion of purity can be simplified: M is pure in N just if

$$N \models (\exists y) ry = m \text{ implies } M \models (\exists y) ry = m$$

for any integer r and any $m \in M$.

4.24 Exercise. Derive the remark in 4.23 from Corollary 4.5.

We denote by $M^{<\kappa}$ the direct sum of κ copies of M and by M^{κ} the direct product of κ copies of M.

4.25 Lemma. *For p.p. formulas without parameters, ϕ and ψ, and any module N, if $\kappa < \omega$ and $n(\phi/\psi; N) < \omega$ then $n(\phi/\psi; N^{\kappa}) = (n(\phi/\psi; N)^{\kappa})$. If $\kappa \geq \omega$, $n(\phi/\psi; N^{\kappa}) = \infty$ or $n(\phi/\psi; N^{\kappa}) = 0$. The analogous results hold for direct sums.*

Proof. For any p.p. formula ϕ, $\phi(N^\kappa) = \phi(N)^\kappa$.

4.26 Corollary. *For any module M and any κ, $M^{<\kappa} \prec M^\kappa$.*

4.27 Historical Notes. The logical analysis of the theory of modules begins with Szmielew's quantifier elimination theorem for Abelian groups ([Szmielew 1955]). The work of Eklof, Fisher, and Sabbagh ([Eklof & Fisher 1972], [Eklof & Sabbagh 1970/71], [Baur 1976]) took a more model theoretic turn. The p.p.-elimination theorem was proved at about the same time but independently by Baur, Garavaglia, and Monk; Garavaglia ([Garavaglia 1979], [Garavaglia 1980]) realized its significance for stability theory. The connections between stability theory and the theory of modules were first pointed out by Shelah ([Shelah 1975]). The exposition here generally follows [Ziegler 1984]. We also rely on [Prest 198?].

5. *Non-Structure Theory*

The bulk of this book is concerned with the development of structure theorems for stable theories. In this section we place that development in context by reciting some of the non-structure theorems and briefly outlining the method of proof. These proofs fall into two parts. The first shows if a theory T is not well behaved then some kind of ordering is definable within the monster model. Second there is a combinatorial argument showing that there are 2^λ of these 'orderings' with power λ for each uncountable λ. In fact, the combinatorial conclusion must be somewhat stronger. The structures constructed must be sufficiently unlike one another that even the Skolem hull of two of the nonisomorphic orderings cannot be isomorphic.

The first result of this kind was Shelah's proof for unstable theories.

5.1 Theorem. *An unstable countable first order theory has 2^λ nonisomorphic models in every uncountable power, λ.*

In this theorem, for example, the required orderings are indeed linear orderings and the first part of the proof depends on showing that every unstable theory contains a formula which linearly orders an infinite set of n-tuples.

A later and stronger result is:

5.2 Theorem. *An unsuperstable countable theory has 2^λ models in every uncountable power.*

This result is proved in Chapter VIII of [Shelah 1978]. The full proof requires a considerable combinatorial arsenal and breaks into cases depending on properties of the cardinal λ. We give a simpler proof for regular λ in Section IX.6 by using more of the forking machinery.

A similar difference occurs between the proof in [Shelah 1982] and the one in [Harrington & Makkai 1985] for showing a theory with the dimensional order property has 2^λ models in every uncountable power. Shelah's proof continues the theme mentioned above of showing an ordering is definable, this time using the language $L_{\omega_1,\omega}$, and then closing under Skolem functions. The more ad hoc argument given here has the advantage that the distinct models can be distinguished in the original language rather than being imposed from without.

5.3 Historical Notes. In addition to Shelah see [Hodges 1983], [Hodges 198?] and [Pillay 1981] for accounts of this technique. The techniques described in this section have been elaborated in several directions. There are two major lines in the study of first order theories. One is the refinement of the sufficient condition to apply the technique. Thus, unstability [Shelah 1971a] is replaced by unsuperstability [Shelah 1975b] and then by the omitting types order property [Shelah 198?]. Alternatively, the conclusion is strengthened to find models which are not only not isomorphic but are mutually non-embeddable. These arguments are of an even more set-theoretic nature [Shelah 1982b], [Shelah 198?d], [Shelah 1983]. In particular, independence results arise in this context [Shelah 1982b].

A further strengthening arises if the non-isomorphic models are required to be elementarily equivalent in various infinitary languages (see for example [Grossberg & Shelah]). Still another generalization considers infinitary conditions in the defining of the trees; here there is considerable work by both Shelah and Grossberg.