Chapter XVI

The Dimensional Order Property

We discuss in this chapter the main dividing line between superstable theories with and without a good structure theory; this dividing line is called, 'the dimensional order property' or just DOP. If a theory has the dimensional order property then it is possible to interpret an arbitrary binary relation into T by considering the dimensions of sets. This 'nonstructure' result leads to the conclusion that T has 2^λ models in every cardinality λ greater than $2^{|T|}$. The great significance of this concept stems from the perhaps even more remarkable consequences of its negation. Essentially, the negation of DOP (NDOP) is the assertion that the relation $p \dashv M$ is a trivial dependence relation. This hypothesis allows one to decompose every model as a tree of small models. There are a number of equivalent formulations of DOP which are useful in various contexts. The name is suggested by the following variant. If T has DOP then there exists a two parameter family $\{p_{\overline{a}\overline{b}}\}$ of copies of a type p such that for any choice of infinite cardinals $\lambda_{\overline{a},\overline{b}}$, there is a model M with $dim(p_{\overline{a}\overline{b}}, M) = \lambda_{\overline{a},\overline{b}}$. This leads to the construction of 2^κ models with cardinality $\kappa \geq 2^{|T|}$ by constructing $\overline{a}_i$, $\overline{b}_i$ for $i < \kappa$ and using $dim(p_{\overline{a}_i\overline{b}_j})$ to encode an arbitrary binary relation on κ.

In the first section of this chapter we discuss the notion of free amalgamation of models in a class K. This leads to the formalization of DOP as an assertion about the triviality of $\dashv$. We develop in Section 2 some technical properties of trivial types which are extremely useful in Section 3 and Chapter XVII. In Section 3, we show that the DOP implies T has many models. In the following chapter we will continue the discussion of theories without the DOP.

1. Avatars of the Dimensional Order Property

We first formally define K-NDOP in a way that will be useful in Chapter XVII and beyond. Then we show that the K in K-NDOP was superfluous; the dimensional order property does not, in fact, depend on the class K. We show that NDOP can be described in terms of finitely generated models and we develop the form of the DOP which makes the name appear most natural.

Throughout this chapter we assume that T is superstable. This is not really necessary and Shelah develops much of the machinery for arbitrary stable theories in [Shelah 1982]. We sacrifice the greater generality for the convenience of a ready supply of regular types. The most immediate advantage of Shelah's added generality is the ability to prove such theorems as, 'A stable but not superstable theory with the dimensional order property has 2^λ models which are λ-saturated and have power λ'. Despite this restriction we often write $\kappa(T)$ when by this convention we mean ω.

Most of the work in Chapter XIV and before was one-dimensional in the sense that we could order our constructions linearly. We now want to consider model constructions where the relations between the building blocks are inherently of a higher dimension. In this book we extend primarily from 1-dimensional diagrams to 2-dimensional diagrams. The full development of classification theory, particuarly the study of infinitary languages [Shelah 1983a], requires the study of n-dimensional diagrams for all finite n.

Much of the work in the remainder of this book is concerned with partially ordered families of sets or, more often, of models. We introduced in Chapter II the notion of an independent system. We denoted the universe of such a system by an upper case Roman letter and the elements with lower case Roman letters, even though it was understood that the elements of the independent system might be subsets of $\mathcal{M}$. In some of the ensuing definitions (specifically a normal family in Section XVII.2) we have to deal with elements of the sets in an independent family. In these more complicated situations we have adopted the convention of referring to the systems as $\bar{A}$, $\bar{B}$, etc., the elements of the system as A, B, etc. and the elements of the elements of the system as a, b, etc. However, we retain the practice of denoting the system by upper case Roman letters and the members (subsets of models) by lower case Roman letters when their elements are not required.

When discussing independent families of models we introduce the following systematic ambiguity.

1.1 Notation. Let $(A, <)$ be a partial order. Let $I \subseteq A$ be an ideal (i.e. downward closed, cf. Definition II.2.24). Then we denote $\bigcup I \subseteq \mathcal{M}$ by I.

It should be clear from context when I denotes a set of subsets of $\mathcal{M}$ and when it denotes the union of that set. Similarly, we use A ambiguously to denote $\cup A$. For example, we will frequently speak of a model M which is prime over (the system) A. This really means M is prime over $\cup A$.

This notation seems to be less cumbersome than that of [Shelah 1982] and [Harrington & Makkai 1985]. They let $\langle \mathcal{I}, < \rangle$ be a partially ordered set and let $\langle A_i : i \in \mathcal{I} \rangle$ be a collection of subsets (of the monster model) indexed by $\mathcal{I}$. For $I \subseteq \mathcal{I}$, A_I denotes $\bigcup_{j \in I} A_j$. Letting subsets of $\mathcal{M}$ act as their own indices allows us to reduce the number of levels of notation by one.

We begin by considering the simplest kind of nonlinear diagram, an independent triangle.

1.2 Definition. i) A triple $\mathcal{N} = \langle N_0, N_1, N_2\rangle$ with $N_1 \downarrow_{N_0} N_2$, N_0 contained in $N_1 \cap N_2$, and each $N_i \in K$ is called a *free K-amalgam.*

ii) A quadruple $\mathcal{N} = \langle N_0, N_1, N_2, N_3\rangle$ such that $\langle N_0, N_1, N_2\rangle$ is a free amalgam, and N_3 is K-prime over $N_1 \cup N_2$ is called a *full (free) K-amalgam.*

Since we deal only with free amalgams we will often omit the adjective free. We say that N_3 *completes* the amalgam $\mathcal{N}$. If the only restriction on the N_i is that they be subsets of the monster model we refer to $\mathcal{N}$ as a set-amalgam.

1.3 Definition. The theory T *does not have the K-dimensional order property,* written T satisfies K-NDOP, if for any full amalgam $\langle N_0, N_1, N_2, N_3\rangle$ of K-models and type p, if $p \dashv N_1$ and $p \dashv N_2$ then $p \dashv N_3$.

If this condition fails we say T has the K-DOP. We show in Theorem 1.7 that we can drop the K and speak just of the *DOP.*

The simplest example of a theory with the dimensional order property is the theory of two crosscutting equivalence relations. (Fig. 1). A model, M, of this theory is easily visualized as a checkerboard. Suppose M is extended to M_1 and M_2 by adding to M a new row for M_1 and a new column for M_2, thus producing a checkerboard with one corner missing. Then the type

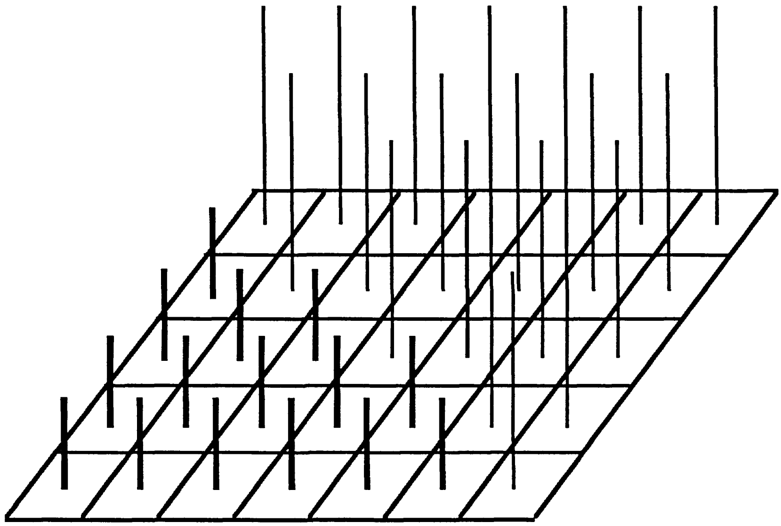

Fig. 1. The dimensional order property (DOP)

of a point in the missing corner square is orthogonal to each of M_1 and M_2 but is realized in every model containing both of them.

The next result sharpens the conditions on DOP and provides a version which is more useful for proving nonstructure results. The three parts of the theorem successively strengthen the result.

1.4 Theorem. *If the full free* **S**-*amalgam* $\mathcal{N}$ *and the type* p *witness the* **S**-*DOP then we may assume*

i) $p \in S(N_3)$ *and* p *is regular.*
ii) N_1 *and* N_2 *are finitely generated over* N_0.
iii) *For* $i = 1, 2$, $N_i = N_0[a_i]$ *where* a_i *realizes a regular type over* N_0.

Proof. i) Since $p \not\perp N_3$, there is a regular $q \in S(N_3)$ with $q \not\perp p$. Write p as $\otimes r_j$ where each r_j is regular. Then, for some j, $r_j \not\perp q$. Now, by the transitivity of nonorthogonality on regular types, if $q \not\perp N_i$ then $r_j \not\perp N_i$ and, a fortiori, $p \not\perp N_i$. Thus we can take q as the desired p.

ii) We assume p satisfies condition i). There is a finite $B \subseteq N_3$ with $|B| < \kappa(T)$ such that p is strongly based on B. Then $t(B; N_1 \cup N_2)$ is **S**-isolated by a finite $A \subseteq N_1 \cup N_2$. Let $\overline{a}_i = A \cap (N_i - N_0)$ for $i = 1, 2$. Now, $\langle N_0, N_0[\overline{a}_1], N_0[\overline{a}_2], N_0[\overline{a}_1, \overline{a}_2]\rangle$ and $p|N_0[\overline{a}_1, \overline{a}_2]$ meet the requirements.

iii) We show that if **S**-NDOP holds for models generated by one element realizing a regular type then it holds for all finitely generated models. Let $\mathcal{N}$ be a free **S**-amalgam with N_1 and N_2 finitely generated over N_0. Let X_1 and X_2 be bases for $R(N_i; N_0)$. The proof is by induction on the sum of the cardinalities of the X_i. Choose $x_i \in X_i$ for $i = 1, 2$ and let $X_i' = X_i - \{x_i\}$. Since $X_1 \cup X_2$ is independent over N_0,

$$\langle N_0[X_1' \cup X_2'], N_0[X_1' \cup X_2], N_0[X_1 \cup X_2'], N_0[X_1 \cup X_2]\rangle$$

is a full **S**-amalgam. Thus, NDOP implies either $p \not\perp N_0[X_1 \cup X_2']$, or $p \not\perp N_0[X_1' \cup X_2]$. Without loss of generality, we may assume the first case occurs. But now, note that $\langle N_0, N_0[X_1], N_0[X_2'], N_0[X_1 \cup X_2']\rangle$ is a full **S**-amalgam so by induction $p \not\perp N_0[X_1]$ or $p \not\perp N_0[X_2']$. In either case we finish, either directly or by applying upward monotonicity for $\not\perp$.

The next lemma helps to explain the name dimensional order property. It shows that in the presence of DOP, if N_3 is **S**-prime over $N_1 \cup N_2$ then N_3 contains an infinite set of indiscernibles over $N_1 \cup N_2$. The dimension of this set of indiscernibles can be fixed arbitrarily to code binary relations into models of T. In particular, consider again the theory of two crosscutting equivalence relations; we are free to place any infinite number of points on each square of the checkerboard.

1.5 Lemma. *Suppose the full amalgam* $\mathcal{N}$ *witnesses that* T *possesses the* **S**-*DOP. Then* N_3 *contains an infinite set of indiscernibles over* $N_1 \cup N_2$.

Proof. (Fig. 2). Let $p \in S(N_3)$ witness the **S**-DOP. Choose $B \subseteq N_3$ with p strongly based on B and $|B| < \kappa(T)$. We will show the following

1.6 Claim. *For any* C *containing* B *with* $|C| < \kappa(T)$ *there is a* $D \supseteq C$ *with* $|D| < \kappa(T)$ *such that* $p|D \vdash p|(N_1 \cup N_2 \cup D)$.

With this fact in mind it is easy to construct the set of indiscernibles. Namely, choose e_i for $i < \omega$ so that e_i realizes $p|(E_i \cup N_1 \cup N_2)$. We can choose the $e_i \in N_3$ since the claim entails that p_i is implied by a type over a set with cardinality $< \kappa(T)$.

We now prove the claim. Using the definition of $\kappa(T)$ and monotonicity, we can extend C to D with $|D| < \kappa(T)$ so that, letting $D_i = D \cap N_i$ for $i < 3$,

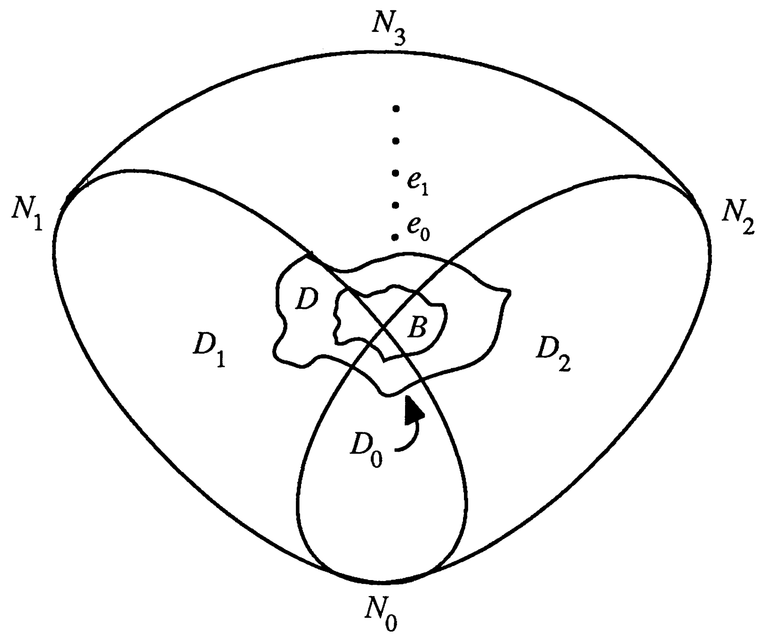

Fig. 2. Lemma XVI.1.5: Simplifying DOP

$$D \underset{D_1 \cup D_2}{\downarrow} (N_1 \cup N_2) \tag{1}$$

and

$$D \underset{D_0}{\downarrow} N_0, D \underset{D_1}{\downarrow} N_1, \text{ and } D \underset{D_2}{\downarrow} N_2. \tag{2}$$

By the definition of a free amalgam we have

$$N_1 \underset{N_0}{\downarrow} N_2. \tag{3}$$

Now, from (1) and the monotonicity properties of forking we can deduce $N_2 \downarrow_{N_1 \cup D_2} N_1 \cup D$. Similarly, (3) implies $N_2 \downarrow_{N_0 \cup D_2} N_1 \cup D_2$. Now, transitivity of independence yields $N_2 \downarrow_{N_0 \cup D_2} N_1 \cup D$. By monotonicity again, we conclude $N_1 \downarrow_{N_0 \cup D} N_2$. The downward preservation of $\dashv$ in the second coordinate implies $p \dashv N_0 \cup D_1$ and $p \dashv N_0 \cup D_2$. Now, using (2) we have $t(N_1; N_0 \cup D) \perp p|(N_0 \cup D)$ and $t(N_2; N_0 \cup D) \perp p|(N_0 \cup D)$. From the strong triviality of orthogonality (Theorem VI.1.19) we conclude $p|(N_0 \cup D) \perp t(N_1 \cup N_2; N_0 \cup D)$. Applying a similar argument to N_0 and D_0 we have $p|D \perp t(N_0; D)$. Since $p|D$ and $p|(N_0 \cup D)$ are stationary $\perp$ implies weak orthogonality (cf. Exercise VI.1.23) so we have $p|D \vdash p|(N_0 \cup D) \vdash p|(N_1 \cup N_2 \cup D)$ as required.

In the light of Theorem X.4.9, Lemma 1.5 asserts that in the presence of DOP prime models over free amalgams are not minimal. The next lemma provides a converse to this result by showing that NDOP implies **S**-prime models over free amalgams are **AT** minimal.

1.7 Lemma. *Let T be a countable ω-stable theory. If T satisfies* **S**-*NDOP and $\mathcal{N}$ is a full* **S**-*amalgam then N_3 is* **AT**-*minimal over $N_1 \cup N_2$. In particular, N_3 is* **AT**-*prime over $N_1 \cup N_2$.*

Proof. (Fig. 3). Suppose for contradiction that there is a proper submodel N_3' with $N_1 \cup N_2 \subseteq N_3' \subseteq N_3$. Then some strongly regular $p \in S(N_3')$ is realized in $N_3 - N_3'$. By NDOP, p is not orthogonal to one of N_1, N_2, say $p \not\dashv N_1$. Applying a consequence of the Three Model Theorem, Theorem XIII.4.3, to N_3, N_3', and N_1, there is a strongly regular type $q \in S(N_3')$ with $q \not\perp p$ such that q does not fork over N_1. Since $N_3 = N_1[N_2]$, for any

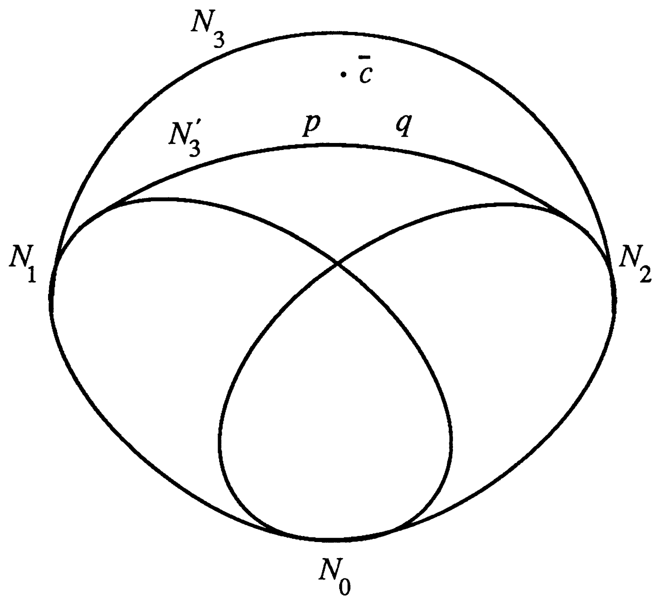

Fig. 3. Lemma XVI.1.7.

$a \in N_3$ $a \not\downarrow_{N_1} N_2$ and a fortiori $a \not\downarrow_{N_0} N_3'$. By Lemma XIII.4.3, $t(N_3; N_3')$ and therefore q is orthogonal to N_0. Since **AT**-strongly regular types are **AT**-minimal some $\bar{c} \in N_3$ realizes q. The choice of q yields $\bar{c} \downarrow_{N_0} N_3'$. Thus, by the definition of $\dashv$, $N_1 \downarrow_{N_0} N_2$ implies $\bar{c} \downarrow_{N_1} N_2$. But as $N_2 \triangleright_{N_1} N_3$, this is impossible and we conclude the theorem.

Corollary XVII.2.4 and the exercise following it give a more satisfying proof of the last result.

The next theorem provides the most intuitive description of the negation of the dimensional order property, the assertion that $\dashv$ is a trivial dependence relation. The proof of the equivalence of this characterization with the original definition is from [Lascar 1985]. One consequence of this result is that we can refer to just the DOP rather than the K-DOP. We discuss another approach to this result in Theorem XVII.1.21.

1.8 Definition. We say $\dashv$ is *trivial* if for any three sets A, B, and C with $B \downarrow_A C$ and any type p, if $p \dashv B$ and $p \dashv C$ then $p \dashv B \cup C$.

The following exercises make the equivalence between NDOP and triviality of $\dashv$ more plausible.

1.9 Exercise. i) Show that $\dashv$ is trivial if and only if for every regular p, if $p \dashv B$ and $p \dashv C$ then $p \dashv B \cup C$.

ii) Suppose $p \dashv M \cup A$ implies $p \dashv M[A]$. Show that for any K, the K-NDOP is equivalent to the assertion that $\dashv$ is trivial.

The hypothesis of Exercise 1.9 ii) does not hold in general. The following more complicated argument of Lascar shows that the conclusion does. We employ the following notation in the proof.

1.10 Notation. Let $\mathcal{N} = \langle N_0, N_1, N_2 \rangle$ be a free K-amalgam. We call the K'-amalgam $\mathcal{N}' = \langle N_0', N_1', N_2' \rangle$ a *parallel amalgam* to $\mathcal{N}$ if the following diagram (Fig. 4) is independent with respect to the partial order indicated by the arrows.

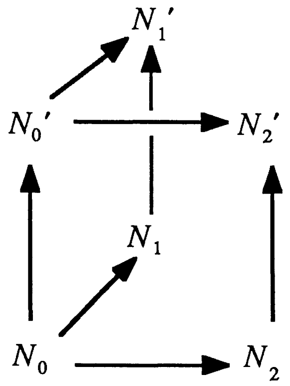

Fig. 4. A parallel amalgam

1.11 Theorem. *The following are equivalent.*

i) *T has the* $\mathbf{S}$*-NDOP.*

ii) *For some acceptable class K, T has the K-NDOP.*

iii) *$\dashv$ is trivial.*

Proof. That i) implies ii) is obvious. To see that ii) implies iii), fix a set-amalgam $\mathcal{N}=\langle N_0, N_1, N_2\rangle$ and assume $p \not\perp N_1 \cup N_2$. Without loss of generality, we may assume p is regular. Extend $\mathcal{N}$ to a parallel amalgam $\langle M_0, M_1, M_2\rangle$ of K-models and let M_3 be K-prime over $M_1 \cup M_2$. Since $p \not\perp N_1 \cup N_2$, $p \not\perp M_3$. By Theorem XIII.3.3, there is a K-strongly regular $q \in S(M_3)$ with $p \not\perp q$. By the K-NDOP, q is not orthogonal to one of M_1 or M_2, say $q \not\perp M_1$. But $M_1 \downarrow_{N_1} N_1 \cup N_2$ implies by Lemma XIII.3.11 that $p \not\perp N_1$ and we finish.

We now show that iii) implies i) (Fig. 5). Let $\mathcal{N} = \langle N_0, N_1, N_2, N_3\rangle$ be a full $\mathbf{S}$-amalgam and suppose $p \in S(N_3)$. We must show $p \not\perp N_1$ or $p \not\perp N_2$. This would be immediate from the definition of triviality but we don't know $p \not\perp N_1 \cup N_2$.

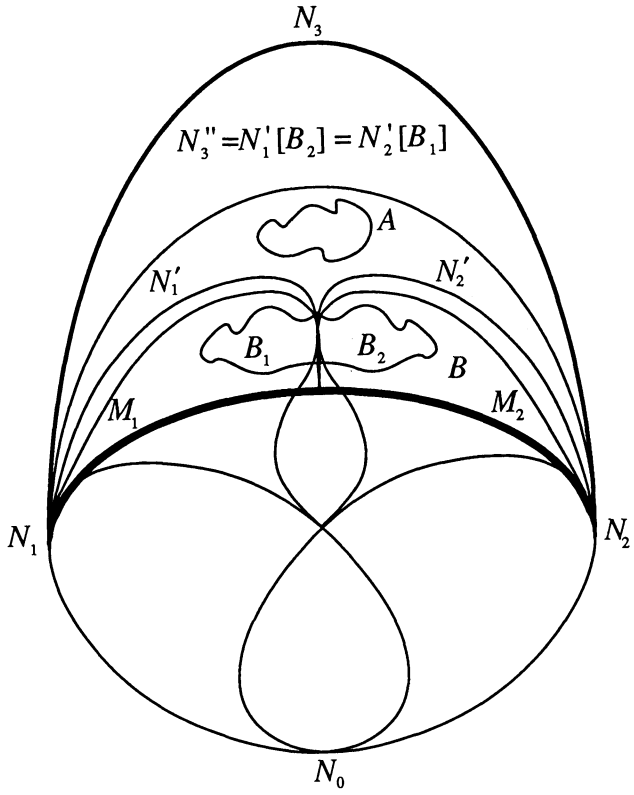

Fig. 5. Theorem XVI.1.11: iii) $\rightarrow$ i)

Let p be strongly based on the finite subset A of N_3. The proof will be by induction on the U-rank of $t(A; N_1 \cup N_2)$. Without loss of generality, p is regular.

If $U(A; N_1 \cup N_2) = 0$ then p does not fork over $N_1 \cup N_2$ so $p \not\dashv N_1 \cup N_2$ and we finish by the triviality of $\dashv$.

Suppose $U(A; N_1 \cup N_2) = \alpha$ and if $\langle M_0, M_1, M_2, M_3\rangle$ is a full amalgam and $p \in S(M_3)$ with p strongly based on the finite subset B of M_3 with $U(B; M_1 \cup M_2) < \alpha$, then $p \not\dashv M_1$ or $p \not\dashv M_2$. Choose N_3' to be **S**-prime over $N_1 \cup N_2$ with $A \downarrow_{N_1 \cup N_2} N_3'$. Let N_3'' be $N_3'[A]$. Let B be a basis for the realizations of regular types over N_3' in N_3''. Since $A \downarrow_{N_1 \cup N_2} N_3'$, for each $b \in B$, $t(b; N_3') \not\perp t(A; N_3')$ which implies $t(b; N_3') \not\dashv N_1 \cup N_2$. Since $\dashv$ is trivial this implies each $t(b; N_3')$ is nonorthogonal to one of N_1 or N_2. By Theorem XIII.3.3 we may without loss of generality replace each $b \in B$ such that $t(b; N_3') \not\dashv N_1$ ($\not\dashv N_2$) by a b' with $b' \downarrow_{N_1} N_3'$ ($b' \downarrow_{N_2} N_3'$). Thus, we can assume B is the disjoint union of B_1 and B_2 where the elements of B_1 (B_2) are independent from N_3' over N_1 (N_2).

By Theorem X.1.28 we can choose copies N_1' and N_2' of $N_3'[B_1]$ and $N_3'[B_2]$ such that $N_3'' = N_1'[B_2] = N_2'[B_1]$. Moreover, we can choose copies M_1 and M_2 of $N_1[B_1]$ and $N_2[B_2]$ respectively so that $N_3'' = N_1'[M_2] = N_2'[M_1]$.

Now let p' be a nonforking extension to $S(N_3'')$ of $p|(A \cup N_1 \cup N_2)$. Note i) $U(A; M_1 \cup N_2') < U(A; N_1 \cup N_2)$ and ii) $M_1 \downarrow_{N_1} N_2'$. The first of these assertions is obvious; the second is a routine application of the forking calculus. Thus we may apply the induction hypothesis to conclude: $p' \not\dashv M_1$ or $p' \not\dashv N_2'$.

In the first case, recall that $N_3 = N_1[N_2]$ and $A \subset N_3$ so $N_2 \downarrow_{N_1} B_1$ implies successively that $N_2 \downarrow_{N_1} M_1$, $N_3 \downarrow_{N_1} M_1$, and $A \downarrow_{N_1} M_1$. But $p \not\dashv A$ and $p \not\dashv M_1$ implies by Theorem XIII.3.11 that $p \not\dashv N_1$ as required.

In the second case we consider the representation of N_3'' as $N_2'[M_1]$ and deduce first that $p \not\dashv M_2$ or $p \not\dashv N_1'$. From the first of these alternatives we deduce as in the previous paragraph that $p \not\dashv N_2$ and finish. If not, we have $p \not\dashv N_1'$ and $p \not\dashv N_2'$. But $N_1' \downarrow_{N_3'} N_2'$ implies, again by Theorem XIII.3.11, that $p \not\dashv N_3'$. But $A \downarrow_{N_1 \cup N_2} N_3'$ yields $p \not\dashv (N_1 \cup N_2)$ and thus by triviality of orthogonality $p \not\dashv N_1$ or $p \not\dashv N_2$ and we finish.

1.12 Exercise. Deduce from the previous theorem that if T is a countable ω-stable theory then T has **S**-NDOP if and only if T has the **AT**-NDOP.

We have shown that the dimensional order property depends only on T and not on a class of models K. If we vary the definition by restricting the type p which can witness the nontriviality of $\dashv$, we do obtain a different notion. The remainder of this section is a first step towards proving Vaught's conjecture for ω-stable theories. To properly count countable models, just as in Chapter XIV, we must distinguish eventually nonisolated types.

1.13 Definition. The theory T *does not have the ENI-dimensional order property*, if for any free amalgam $\langle N_0, N_1, N_2, N_3\rangle$ of **S**-models and ENI type p, if $p \dashv N_1$ and $p \dashv N_2$ then $p \dashv N_3$. We write T satisfies ENI-NDOP.

Naturally, the negation of this property is referred to as the ENI-DOP. Notice that the theory CER_2 of two cross-cutting equivalence relations is an $\aleph_0$-categorical, ω-stable theory with DOP but which does not have ENI-DOP.

1.14 Exercise. Give an example of an ω-stable but not $\aleph_0$-categorical theory with DOP but without ENI-DOP.

1.15 Exercise. Show that if T has the ENI-DOP then T has the DOP.

Since we showed in Section XV.2 that nonorthogonality preserved ENI types, the following result can be deduced immediately from Theorem 1.4.

1.16 Theorem. *If the full* $\mathbf{S}$*-amalgam* $\mathcal{N}$ *and the ENI-type* p *witness the ENI-DOP then we may assume*

i) $p \in S(N_3)$.
ii) N_1 *and* N_2 *are finitely generated over* N_0.
iii) *For* $i = 1, 2$, $N_i = N_0[a_i]$ *where* $t(a_i; N_0)$ *is strongly regular.*

1.17 Historical Notes. The dimensional order property was first defined and applied in [Shelah 1982]. Our treatment depends greatly on later expositions by [Harrington & Makkai 1985] and [Lascar 1985]. In addition to proving that NDOP is equivalent to the triviality of $\dashv$, Lascar introduces in that paper another interesting notion. He says a type p is bounded by a set A if when writing p as $\otimes r_i$, where each r_i is regular, all of the r_i are not orthogonal to A. Lascar calls a theory without the dimensional order property presentable. This terminology is justified by the decomposition theorem in Chapter XVII.

2. Triviality of Forking

In this section we develop some technical consequences of the assumption that forking is trivial on the realizations of a type p. In Section XV.3 we showed that if there exists a type $p \in S(\overline{a})$, p is orthogonal to the empty set and $t(\overline{a}; \emptyset)$ is a nontrivial weight one type then there exist many non-isomorphic models. In this section we expound the properties of trivial types. We will rely on the results here in the proof that DOP implies a theory has the maximal number of models and in the proof of the Vaught and Morley conjectures for ω-stable countable theories.

Recall from Section XV.2 the definition of a trivial type.

2.1 Definition. The stationary type $p \in S(A)$ is *trivial* if for any nonforking extension p' of p, any three pairwise independent realizations of p' are in fact independent.

We call a triple of points which form a counterexample to triviality a *triangle*. Clearly, if p is trivial so is any nonforking extension of p. Given a triple $\langle \bar{a}, \bar{b}, \bar{c} \rangle$ of realizations of $p \in S(A)$ and $B \supseteq A$, we can choose $\langle \bar{a}', \bar{b}', \bar{c}' \rangle$ realizing $t(\bar{a}^\frown \bar{b}^\frown \bar{c}; A)$ with $\bar{a}'^\frown \bar{b}'^\frown \bar{c}' \downarrow_A B$. Since the first triple is a triangle over A if and only if the second is a triangle over B we deduce the invariance of triviality under nonforking extensions.

2.2 Proposition. *If p' is a nonforking extension of the stationary type p, then p is trivial if and only if p' is trivial.*

Using this fact, it is easy to prove the following proposition.

2.3 Proposition. *If p is a trivial type and I is a set of pairwise independent realizations of p then I is independent.*

When p is regular we can extend the triviality property to sets which do not realize p.

2.4 Lemma. *Let $p \in S(A)$ be trivial and regular. If $\bar{a}$ realizes p and for some B_1, B_2 with $B_1 \downarrow_A B_2$, $\bar{a} \not\downarrow_A B_1 \cup B_2$ then $\bar{a} \not\downarrow_A B_1$ or $\bar{a} \not\downarrow_A B_2$.*

Proof. Without loss of generality we can replace A by a strongly $\kappa(T)$-saturated model M. Thus, p denotes $t(\bar{a}; M)$. For $i = 1, 2$, let $D_i = p(M[B_i])$. Then, $p \perp t(M[B_i]; M[D_i])$. For, if not, p would be realized in $M[B_i] - M[D_i]$ contradicting the choice of the D_i. Now since $M[B_1] \downarrow_M M[B_2]$, $M[B_1] \downarrow_{M[D_1] \cup M[D_2]} M[B_2]$, and for each i $t(M[B_i]; M[D_1] \cup M[D_2]) \parallel t(M[B_i]; M[D_i])$. Since orthogonality is preserved by parallelism and is trivial we have $t(M[B_1] \cup M[B_2]; M[D_1] \cup M[D_2]) \perp p$. Now, if $\bar{a} \downarrow_M (D_1 \cup D_2)$, then $\bar{a} \downarrow_M M[D_1] \cup M[D_2]$. By the definition of orthogonality we conclude

$$\bar{a} \underset{M[D_1] \cup M[D_2]}{\downarrow} M[B_1] \cup M[B_2].$$

Thus, $\bar{a} \not\downarrow_M B_1 \cup B_2$ implies $\bar{a} \not\downarrow_M D_1 \cup D_2$. So by triviality of p, $\bar{a} \not\downarrow_M D_1$ or $\bar{a} \not\downarrow_M D_2$. But the first implies $\bar{a} \not\downarrow_M B_1$ and the second implies $\bar{a} \not\downarrow_M B_2$ so we finish.

The following corollary is almost immediate.

2.5 Corollary. *Suppose that I is a set of realizations of the trivial regular type $p \in S(A)$ and that I is independent over A. If each $a \in I$ satisfies $a \downarrow_A B$ then I is an independent set over B.*

Proof. Since p is trivial, it suffices to show that I is pairwise independent over B. Thus, it is more than enough to show that if $a, b \in I$ then $a \downarrow_A B \cup b$. But by Lemma 2.4, the last assertion follows from $a \downarrow_A B$, $b \downarrow_A B$, and $a \downarrow_A b$, which all hold.

We would like to extend Lemma 2.4 from trivial regular types to trivial weight one types. In fact, we will have to strengthen the hypothesis on $B \cup C$. Before proving the extension we must show that triviality is a property of the nonorthogonality class of a weight one type.

2.6 Theorem. *If p and q are nonorthogonal weight one stationary types and q is trivial then p is trivial.*

Proof. Without loss of generality, assume that p and q are in $S(M)$ for some strongly $\kappa(T)$-saturated model M. Since p and q are nonorthogonal, they are not almost orthogonal. Thus we can define a map f from $p(\mathcal{M})$ into $q(\mathcal{M})$ such that $a \not\downarrow_M f(a)$. Since q has weight one, if a_1 and a_2 are independent so are $f(a_1)$ and a_2. Now we can see that f preserves pairwise independence. For if a_1 and a_2 were independent but $f(a_1)$ and $f(a_2)$ were not, the assumption that $t(f(a_2); M)$ has weight one would be contradicted. Now, suppose that $A = \{a_1, a_2, a_3\}$ are pairwise independent realizations of p. Note that f is 1-1 on A. Then f^{-1} satisfies the conditions of Theorem XIII.2.9 so A is independent.

Another easy consequence of Lemma 2.4 strengthens the triviality of forking to include all realizations of trivial regular types.

2.7 Corollary. *If p and q are trivial regular types then forking is trivial on $p(\mathcal{M}) \cup q(\mathcal{M})$.*

Now we can obtain the promised extension of Lemma 2.4 to weight one types. The following proof relies on superstability by assuming the existence of regular types. I don't know whether this assumption can be avoided.

2.8 Theorem. *If $a \not\downarrow_M c_1{}^\frown c_2$ with $c_1 \downarrow_M c_2$, $t(a; M)$ is trivial, and the type of each of a, c_1, c_2 over M has weight one then $a \not\downarrow_M c_1$ or $a \not\downarrow_M c_2$.*

Proof. Without loss of generality, we assume that M is strongly $\kappa(T)$-saturated. Since $p \not\perp t(c_1{}^\frown c_2; M)$, p is realized by some $a' \in M[c_1, c_2]$. Form $M[a'] \prec M[c_1, c_2]$ and choose $d \in M[a']$ to realize a regular type q over M. Then $p \not\perp q$ so by Lemma 2.6, q is trivial. Since $d \in M[c_1, c_2]$, $d \not\downarrow_M c_1{}^\frown c_2$. Applying Lemma 2.4, we can assume without loss of generality that d depends on c_1 over M. Since q has weight one, d and c_2 are independent. If d depends on a over M, transitivity of forking on weight one types yields the result. But if a and d are independent, applying Lemma 2.4 again, we have $d \downarrow_M a{}^\frown c_2$. Thus, $d \downarrow_{M \cup c_2} a$. But $a \not\downarrow_M c_1{}^\frown c_2$ and $a \downarrow_M c_2$ implies $a \not\downarrow_{M \cup c_2} c_1$. Similarly, $d \not\downarrow_M c_1{}^\frown c_2$ and $d \downarrow_M c_2$ implies $d \not\downarrow_{M \cup c_2} c_1$. The last three assertions contradict the hypothesis that $t(c_1; M \cup c_2)$ has weight one so we finish.

The following two exercises show the true strength of Theorem 2.8. The first one shows that if an element a depends on an independent set I of realizations of a type then over an appropriate base it depends on a subset I_0 with $|I_0| \leq 2$. Combining this observation with Theorem 2.8 yields Exercise 2.10.

2.9 Exercise. Suppose p and q are nonalgebraic stationary types over A. Show that if a realizes p and I is a minimal set of realizations of q such that $a \not\downarrow_A I$ and $|I| \geq 2$ then there is an S-saturated model M and realizations c_1, c_2 of q such that $a \not\downarrow_M c_1{}^\frown c_2$ but $a \downarrow_M c_1$, $a \downarrow_M c_2$.

2.10 Exercise. Suppose I and J are independent sequences realizing trivial weight one types $p, q \in S(A)$ such that each $\overline{a} \in I$ satisfies $\overline{a} \not\downarrow_A J$ and each $\bar{b} \in J$ satisfies $\bar{b} \not\downarrow_A I$. Show $\bar{a} \not\downarrow_A \bar{b}$ establishes a 1-1 correspondence between I and J.

Another version of this result continues the program of finding sufficient conditions to make orthogonality and almost orthogonality equivalent. We began by proving the equivalence if the two types are over a strongly saturated model (Theorem VI.1.40); later we extended the result to arbitrary models of a countable ω-stable theory (Corollary XIII.3.8.) Here we move in another direction by imposing no condition on the set but requiring that the types be trivial. [Buechler 1986], extending [Cherlin, Harrington, & Lachlan 1985], shows that trivial can be replaced by modular. We will rely heavily on the next two results in Chapter XVIII.

2.11 Corollary. *If $p, q \in S(A)$ are stationary trivial weight one types then $p \not\perp q$ implies $p \not\perp^a q$.*

Proof. By Corollary VI.2.18, $p \not\perp q$ implies there are finite independent sequences E, F of realizations of p, q, respectively, so that $t(E; A) \not\perp^a t(F; A)$. But since p and q are trivial, Theorem 2.9 yields that for some $e \in E$, $f \in F$, $e \not\downarrow_A f$ and we finish.

2.12 Lemma. *Let $p, q \in S(A)$ be stationary trivial types with $p \not\perp q$. Suppose p is regular and for some ϕ, (q, ϕ) is strongly regular. Then for every b realizing p and every $M \supseteq A \cup b$, there is a $c \in M$ which realizes q and $b \not\downarrow_A c$.*

Proof. By Theorem 2.11, $p \not\perp^a q$. Let b realize p and choose d realizing q with $b \not\downarrow_A d$. Then there is a formula $\psi(x, y)$ such that for any e, $\models \psi(b, e)$ implies $b \not\downarrow_A e$. Fix $M \supseteq A \cup b$ and choose $c \in M$ so that $\models \phi(c) \wedge \psi(b; c)$. Since (q, ϕ) is strongly regular either c realizes q and we finish or $t(c; A) \perp q$. But $t(c; A) \not\perp p$ and $p \not\perp q$ so, since p is regular, $t(c; A) \not\perp q$.

The important Theorems 2.9 and 2.11 require only that p and q have weight one. The next attribute holds for regular trivial types but not for general weight one trivial types.

2.13 Definition. The type $p \in S(A)$ is *totally trivial* if for any $\langle \bar{a}, \bar{b}, \bar{c} \rangle$ realizing p, if $\bar{a} \not\downarrow_A \bar{b}^\frown \bar{c}$ then $\bar{a} \not\downarrow_A \bar{c}$ or $\bar{a} \not\downarrow_A \bar{b}$.

Using the transitivity characterization of regularity, it is easy to see

2.14 Proposition. *Every regular trivial type is totally trivial.*

The next example shows the necessity of regularity for the last proposition.

2.15 Example. (Fig. 6). Let the language of T contain unary predicates R (for regular) and W (for weight one), a binary relation E, and a ternary relation S. For simplicity, in the following description we identify each relation symbol of L with its interpretation in a model of T. Roughly, a model of T consists of two sets R and W. There is an equivalence relation with infinitely many infinite classes on W. Each class is named by an element of R. There is a graph defined on the elements of W. Each component of the graph lies in a single equivalence class. The graph is symmetric and contains no cycles. Finally, each edge is labeled by an element of R. For each element a of W and each element b of R, exactly one other element of W is connected to a by an edge labeled by b.

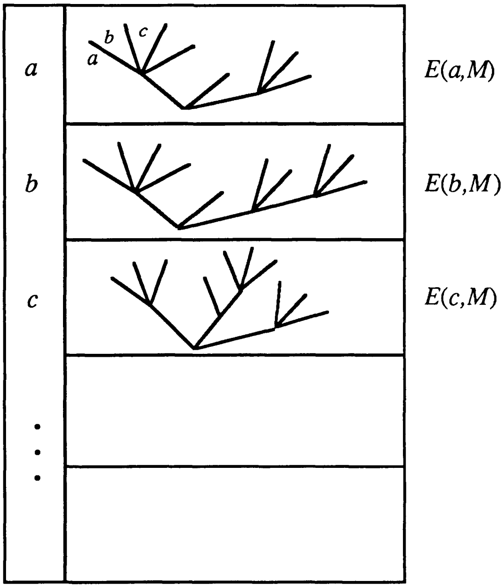

Fig. 6. Example XVI.2.15.

More formally, $E \subseteq R \times W$; each element of R is connected via E with infinitely many elements of W but each element of W is connected to exactly one element of R. Moreover, $S \subseteq R \times W \times W$. If $\langle r, w_1, w_2 \rangle \in S$ then $\langle r, w_2, w_1 \rangle \in S$ and for some r' both $\langle r', w_1 \rangle$ and $\langle r', w_2 \rangle$ are in E. Thus, the projection of S on $W \times W$ defines a symmetric graph on W which has no cycles and connects only points in the same equivalence class defined by R and E. The first coordinate of S determines a labeling of the edges of this graph by elements of R. Namely, each point will be related to infinitely many other points by the relation $(\exists x)R(x) \wedge S(x, y, z)$. There is a $1-1$ correspondence between the points in R and the points related to an element a. That is, for each $a \in W$, f_a is a $1-1$ map from R to the neighbors of a in the graph given by $f_a(x)$ is the unique z such that $S(x, a, z)$.

Now we claim, first, T is ω-stable. To verify this one must first check that (after adding predicates which indicate the distance between pairs of points in the graph and the possible equalities among the labels of the edges of a path connecting the pair) T admits elimination of quantifiers and then count types. Second, each 1-type containing $R(x)$ is regular and trivial. Here are two key observations to support this claim. Any permutation of R which preserves the number of components in the graph associated with a point of R extends to an automorphism of the model. Any permutation of an equivalence class which is a homomorphism for the expanded language extends to an automorphism of the model.

But, some 1-types which contain $W(x)$ have weight one and are trivial but are not totally trivial. If $W(x) \in q \in S(M)$ then the Morley rank (or U-rank) of q is determined by the 'distance' of a realization of q from M. It is easy to check that the rank of q is at most $\omega + 1$ and this rank is attained by the type $\hat{q}$ which asserts that x is not related by E to any element of M. Now, if $a, b \in W$ realize $\hat{q}$, they are independent over M if and only if they are in distinct classes of the partition. Thus, forking is trivial on $\hat{q}$. But, suppose a and b are in the same class of the partition and there exists a $c \in R$ with $S(c, a, b)$ and $\neg E(c, a)$. Then for any $d \in W$ with $E(c, d)$, d is independent from each of a and b but depends on $a^\frown b$.

The following lemma connects the dimensional order property with the existence of non-trivial types. For convenience, we say a type r has depth at least one if there is a nonforking extension p of r with the following properties. For some **S**-model M $p \in S(M)$, and there is a realization $\overline{c}$ of p, and a regular type $q \in S(M[\overline{c}])$ with $q \dashv M$. This is a special case of the definition of the depth of a type in Chapter XVII.

2.16 Lemma. *If T does not have the dimensional order property then T has no nontrivial regular type p with depth one or more.*

Proof. Without loss of generality assume $p \in S(M)$ and M is S-saturated. Suppose $\langle \overline{a}, \overline{b}, \overline{c} \rangle$ form a triangle which witnesses the nontriviality of p and that p has depth at least one. Choose $A \subseteq M$ with $|A| < \kappa(T)$ such that $t(\overline{a}^\frown \overline{b}^\frown \overline{c}; M)$ is strongly based on A. Form $M[\overline{a}, \overline{b}]$. Now $t(\overline{c}; A \cup \overline{a}^\frown \overline{b})$ is realized in $M[\overline{a}, \overline{b}]$ by some $\overline{c}'$. If $\overline{c}' \not\downarrow_A M$ then $t(\overline{c}'; M) \perp p$. But since $\overline{a}$ and $\overline{b}$ are independent realizations of p, the only regular types realized in $M[\overline{a}, \overline{b}]$ are those which are not orthogonal to p. Thus, $\overline{c}' \downarrow_A M$. Suppose for contradiction that $\overline{c}' \not\downarrow_M \overline{a}$. Then $t(\overline{c}'; M \cup \overline{a}) \perp p$. Hence $\overline{c}' \downarrow_{M \cup \overline{a}} \overline{b}$. But, since $\overline{b} \downarrow_A M \cup \overline{a}$, transitivity of independence yields $\overline{c}' \downarrow_A \overline{b}$. This contradicts the choice of $\overline{c}'$. Now, since $\overline{a} \downarrow_M \overline{c}'$ and $\overline{b} \downarrow_M \overline{c}'$, $p_{\overline{c}'}$ is orthogonal to both $M[\overline{a}]$ and $M[\overline{b}]$ which gives an example of the DOP.

2.17 Historical Notes. This section appeared in [Baldwin & Harrington]. Lemma 2.16 is from [Saffe 1983]. Many of the other results were implicit in the main gap papers [Shelah 1982]. We mentioned after Lemma V.2.6 and in Section X.4 the similarity between an infinite set of indiscernibles and a model. The same phenomenon occurs in [Shelah 1986c] where Shelah proves a version of Lemma 2.12 replacing M by the algebraic closure of an infinite set of indiscernibles.

3. DOP Implies Many Nonisomorphic Models

In this section we prove any superstable theory with DOP has 2^λ models in every $\lambda \geq 2^{|T|}$ and an ω-stable countable theory with ENI-DOP has $2^{\aleph_0}$ countable models. To obtain a precise common statement of these two results, recall from Section XIV.2 that a type p over a set A is (μ, K)-tractable if there are fewer than μ realizations of p in a K-prime model over A.

3.1 Theorem. *Let T be a superstable theory with the DOP. Suppose the (λ, K)-tractable type p witnesses the DOP and $\lambda \geq \lambda_0(\mathbf{I})$. If $\mu \geq \lambda_0(\mathbf{I})$ then $I(\mu, K) = 2^\mu$.*

Remember that for $K = \mathbf{S}$, $\lambda_0(\mathbf{S}) = 2^{|T|}$ while for countable ω-stable T, $\lambda_0(\mathbf{AT}) = \aleph_0$. Our argument follows that of [Harrington & Makkai 1985]. Shelah's original proof [Shelah 1982] was more along the line discussed in Section I.5.

We begin by translating the DOP into a collection of data about a finite set of finite sequences. We will use these finite sequences to code a family of graphs into models of T with power λ and thus deduce the theorem.

The following reduction depends on several ideas. The first is that because of superstability, we can reduce the relations between the models which witness DOP according to Definition 1.3 to relations between finite sequences. The second is that because N_1 and N_2 were K-minimal (Theorem 1.4), we can demand that these finite sequences have weight one. Finally, since N_1 and N_2 were prime over K-strongly regular types we can demand either that $(N_1 \perp N_2; N_0)$ or, invoking Corollary XII.1.15, that N_1 is isomorphic to N_2 over N_0.

To apply the following lemma to a countable ω-stable theory recall that for such a theory a nonprincipal type is $(\aleph_0, \mathbf{AT})$-tractable.

3.2 Reduction. Suppose T is superstable and a (λ, K)-tractable type witnesses that T has DOP. Then, there exist finite sequences $\overline{n}_0$, $\overline{n}_1$, $\overline{n}_2$, and $\overline{n}_3$ such that $\overline{n}_0 \subseteq \overline{n}_1 \subseteq \overline{n}_3$, $\overline{n}_0 \subseteq \overline{n}_2 \subseteq \overline{n}_3$,

$$\overline{n}_1 \underset{\overline{n}_0}{\downarrow} \overline{n}_2,$$

$$\overline{n}_1 \rhd_{\overline{n}_2} \overline{n}_3,$$

$$\overline{n}_2 \rhd_{\overline{n}_1} \overline{n}_3,$$

and

$$\overline{n}_1 \cup \overline{n}_2 \rhd_{\overline{n}_0} \overline{n}_3.$$

Furthermore, there is a (λ, K)-tractable type $p \in S(\overline{n}_3)$ with $p \dashv \overline{n}_1$ and $p \dashv \overline{n}_2$. Finally, if $q_1 = t(\overline{n}_1; \overline{n}_0)$ and $q_2 = t(\overline{n}_2; \overline{n}_0)$ we can guarantee that q_1 and q_2 are stationary weight one types and either $q_1 = q_2$ or $q_1 \perp q_2$.

Proof. (Fig. 7). By Theorem 1.4 (or Theorem 1.16 for the ω-stable case) we can find a full amalgam of K-models, $\langle N_0, N_1, N_2, N_3 \rangle$ and a (λ, K)-

tractable type $r \in S(N_3)$ such that $r \dashv N_1$, $r \dashv N_2$. Moreover, we can insist that for $i = 1, 2$, $N_i = N_0[\overline{a}_i]$ where $t(\overline{a}_i; N_0)$ is K-strongly regular. If $t(\overline{a}_1; N_0)$ and $t(\overline{a}_2; N_0)$ are not orthogonal, Corollary XII.1.15 implies that we can find an isomorphism which fixes N_0, takes $\overline{a}_1$ to $\overline{a}_2$ and N_1 to N_2.

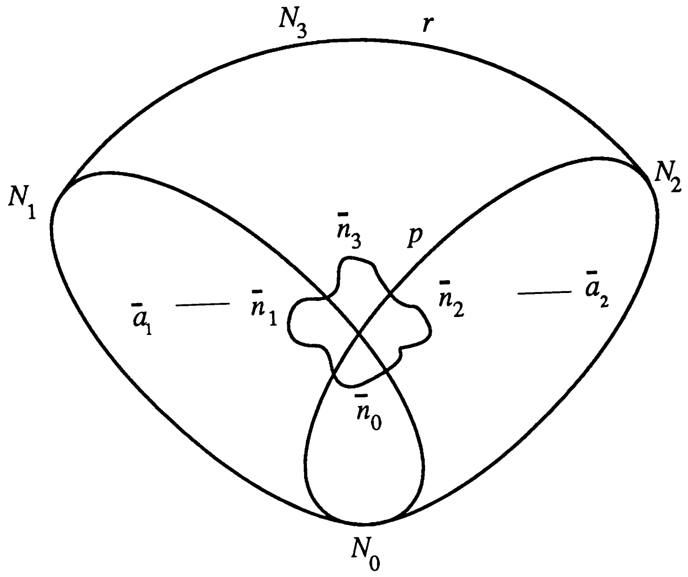

Fig. 7. Reducing to finite sets

Choose $\overline{n}_3 \subseteq N_3$ so that r is strongly based on $\overline{n}_3$. Let $p = r|\overline{n}_3$. Then choose $\overline{n}_i$, for $i = 1, 2$, so that $t(\overline{n}_3; N_i)$ is strongly based on $\overline{n}_i$. If $t(\overline{a}_1; N_0) \perp t(\overline{a}_2; N_0)$, clearly $t(\overline{n}_1; N_0) \perp t(\overline{n}_2; N_0)$. If not, replace $\overline{n}_2$ by $\overline{n}_2$ union the image of $\overline{n}_1$ under the isomorphism mapping N_1 to N_2 and similarly for $\overline{n}_1$. This replacement guarantees that $q_1 = q_2$. Now choose $\overline{n}_0 \subseteq N_0$ so that $t(\overline{n}_1{}^\frown\overline{n}_2; N_0)$ is strongly based on $\overline{n}_0$. Since for $i = 1, 2$, $N_i = N_0[\overline{a}_i]$ and $t(\overline{a}_i; N_0)$ is K-strongly regular, each $t(\overline{n}_i; N_0)$ has weight one. Since N_3 is K-prime over $N_1 \cup N_2$, $\overline{n}_1 \rhd_{\overline{n}_2} \overline{n}_3$ and $\overline{n}_2 \rhd_{\overline{n}_1} \overline{n}_3$. Similarly, $\overline{n}_1 \cup \overline{n}_2 \rhd_{\overline{n}_0} \overline{n}_3$.

We can make the following simplifications in notation.

3.3 Notation. We will replace the vector notation by using individual letters, a, b, c to refer to finite tuples. Specifically we denote $\overline{n}_i$ as n_i for i=0,1,2,3. We will assume that the sequence $\overline{n}_0$ is named in the language. In particular this means that q_1, q_2 are stationary types over the empty set. For any model M we denote by $Q(M)$ the set of all tuples which realize either q_1 or q_2.

We are going to build many models by coding graphs on realizations of q_1 and q_2. The next result will simplify picking out the graph in the model. It generalizes Corollary 2.7 (since the q_i are only weight one rather than regular). We are able to prove the stronger result because of the additional information that the q_i are either equal or orthogonal.

3.4 Proposition. i) *q_1 and q_2 are trivial types.*
ii) *Forking is trivial on $Q(\mathcal{M})$.*

Proof. i) We have that $n_1 \rhd_{n_2} n_3$. Since $p \dashv n_1$ and $wt(t(n_1; n_2)) = 1$, by Theorem XV.2.12, either we are finished or $t(n_1; n_2)$ is trivial. Since $n_1 \downarrow_{n_0} n_2$, this implies q_1 is trivial. The argument for q_2 is entirely symmetric.

ii) If $q_1 = q_2$ then ii) is just a restatement of i). So suppose $q_1 \perp q_2$ and $I \subseteq Q(\mathcal{M})$ is pairwise independent but not independent and $|I|$ is minimal among all such sets. Then we can write I as $C \cup D \cup \{b\}$ where $C \subseteq q_1(\mathcal{M})$ and $D \cup \{b\} \subseteq q_2(\mathcal{M})$. By the triviality of q_2, $b \downarrow D$ and by orthogonality of q_1 and q_2, $b \downarrow C$. To get $b \downarrow C \cup D$, it suffices by Lemma II.2.5 to show $b \downarrow_D C$. This holds because $t(b; \emptyset) \perp t(C; \emptyset)$, $C \downarrow D$, and by induction $b \downarrow D$.

To reduce eyestrain at the possible cost of straining the reader's short term memory, we introduce a number of notions which are used only in this proof. Here is the first set.

3.5 Notation. A *germane pair* $\langle a, b\rangle$ is a realization of $t(n_1, n_2; \emptyset)$; a *germane triple*, $\langle a, b, e_{ab}\rangle$ is a realization of $t(n_1, n_2, n_3; \emptyset)$. If $\langle a, b, e_{ab}\rangle$ is a germane triple, we denote by p_{ab} a copy of $p \in S(e_{ab})$ which is orthogonal to both a and b.

Note that if $\langle a, b\rangle$ is a germane pair then $a \downarrow b$. Although we write e_{ab} and p_{ab}, there are many choices for this set and type for any given a, b.

3.6 Exercise. If $\langle a, b\rangle$ is a germane pair and $q_1 = q_2$ then $t(a^\frown b; \emptyset) = t(b^\frown a; \emptyset)$.

The principal technical difficulty of this proof centers around the following problem. Given p_{ab} and p_{cd} with $p_{ab} \not\perp p_{cd}$; show $a \not\downarrow c$ and $b \not\downarrow d$. This does not hold in general; the next few lemmas lay out special conditions where it does hold. We will then arrange the construction so we can obtain those conditions.

First suppose that $b = d$ or $a = c$.

3.7 Proposition. i) *If $\langle a, b\rangle$ and $\langle a', b\rangle$ are germane pairs and $p_{ab} \not\perp p_{a'b}$ then $a \not\downarrow a'$.*
ii) *If $\langle a, b\rangle$ and $\langle a, b'\rangle$ are germane pairs and $p_{ab} \not\perp p_{ab'}$ then $b \not\downarrow b'$.*

Proof.

i) If $a \downarrow a'$, since forking is trivial on $Q(\mathcal{M})$, $\{a, a', b\}$ is an independent set. As $a \rhd_b e_{ab}$ and $a' \rhd_b e_{a'b}$, we can conclude $e_{ab} \downarrow_b e_{a'b}$. Since $p_{ab} \dashv b$, this implies $p_{ab} \perp p_{a'b}$.
ii) If $q_1 = q_2$, ii) is a restatement of i). If not, it still follows by a symmetric proof.

To consider the case when a, b, c, d are distinct we introduce some further notation.

3.8 Notation. A 4-tuple $\langle a, b, c, d\rangle$ is in *normal position* if $\langle a, b\rangle$ and $\langle c, d\rangle$ are germane pairs and $a \downarrow d$ and $b \downarrow c$.

Note that $q_1 \perp q_2$ implies that any two germane pairs are in normal position. The notion of normal position can be visualized by regarding a, b, c, d as the vertices of a square. We will connect two vertices of the square to indicate the points are independent. The germane pairs guarantee the top and bottom are independent (draw the top and bottom). (Fig. 8). The additional requirement of normal position is that the diagonals are independent (draw the diagonals). (Fig. 9). Thus, if $\langle a, b, c, d\rangle$ is in normal position, a and c realize q_1 while b and d realize q_2.

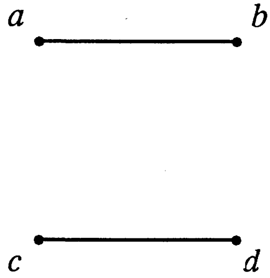

Fig. 8. Germane pairs

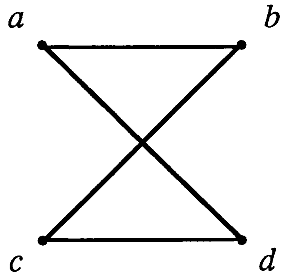

Fig. 9. Normal position

3.9 Exercise. If $\langle a, b, c, d\rangle$ is in normal position and $q_1 = q_2$ then $d \downarrow a^\frown b$.

The next proposition concerns whether we can draw the sides.

3.10 Proposition. *If $\langle a, b, c, d\rangle$ is in normal position with $a \downarrow c$ and $b \downarrow d$ then $p_{ab} \perp p_{cd}$.*

Proof. (Fig. 10). The hypotheses and the triviality of forking on $Q(\mathcal{M})$ yield that $\{a, b, c, d\}$ is an independent set and, in particular, $ab \downarrow cd$. Since $ab \rhd e_{ab}$ and $cd \rhd e_{cd}$, we deduce using Theorem VI.2.21 from $p_{ab} \dashv \emptyset$ that $p_{ab} \dashv e_{cd}$ and thus $p_{ab} \perp p_{cd}$.

Note that the argument for Proposition 3.10 depends only on the properties of the set $\{a, b, c, d\}$, not the ordered pair $\langle a, b\rangle$. Thus, we have

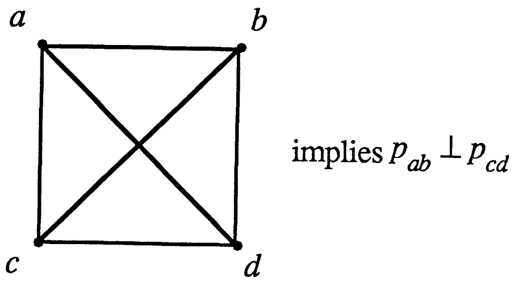

Fig. 10. Proposition XVI.3.10.

3.11 Exercise. Assume $\langle a, b, c, d\rangle$ is in normal position. Then $p_{ba} \perp p_{cd}$, $p_{ba} \perp p_{dc}$, $p_{ab} \perp p_{cd}$, and $p_{ab} \perp p_{dc}$.

This shows if $\langle a, b, c, d\rangle$ is in normal position and $p_{ab} \not\perp p_{cd}$ then at least one of $a \not\downarrow c$ and $b \not\downarrow d$ holds. We want them both to hold; this requires some further hypotheses.

3.12 Proposition. (Fig. 11). *Suppose $\langle a, b, c, d\rangle$ is in normal position and $p_{ab} \not\perp p_{cd}$.*

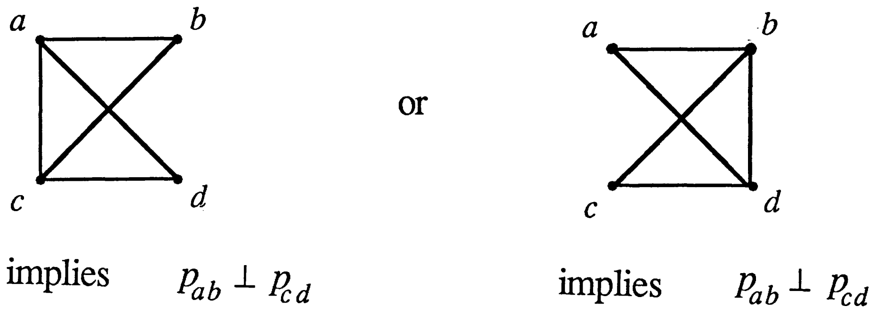

Fig. 11. Proposition XVI.3.12.

i) *If $q_1 = q_2$ then $a \not\downarrow c$ and $b \not\downarrow d$.*
ii) *If $q_1 \perp q_2$ and $ac \downarrow bd$ then $a \not\downarrow c$ and $b \not\downarrow d$.*

Proof. i) By Proposition 3.10, we know $a \not\downarrow c$ or $b \not\downarrow d$. Suppose $a \not\downarrow c$ and $b \downarrow d$. Since $q_1 = q_2$, $t(a^\frown b; \emptyset) = t(c^\frown d; \emptyset)$. Let α be an automorphism which interchanges a and b. Choose $\langle \hat{c}, \hat{d}\rangle$ to realize $t(\alpha(c), \alpha(d); a^\frown b)$ and with $\hat{c}\hat{d} \downarrow_{ab} cd$. We have $p_{cd} \not\perp p_{ab}$ and $p_{ab} \not\perp p_{\hat{d}\hat{c}}$ so $p_{cd} \not\perp p_{\hat{d}\hat{c}}$. As germane pairs, we have $c \downarrow d$ and $\hat{c} \downarrow \hat{d}$. We chose $\hat{d}\hat{c} \downarrow_{ab} cd$ and were given $d \downarrow ab$ so $d \downarrow \hat{d}$ and $d \downarrow \hat{c}$. Now $c \downarrow b$, $b = \alpha(a)$, and $\hat{c}$ realizes $t(\alpha(c); b)$ implies $\hat{c} \downarrow a$. Similarly,

$\hat{d} \downarrow a$. But $c \not\downarrow a$ so since q_1 has weight one $c \downarrow \hat{c}$ and $c \downarrow \hat{d}$. By Exercise 3.11 we conclude $p_{cd} \perp p_{\hat{d}\hat{c}}$ and this contradiction yields part i).

ii) Suppose for contradiction that $b \downarrow d$. By monotonicity from $ac \downarrow bd$ we conclude

$$c \underset{a}{\downarrow} b \tag{1}$$

and

$$a \underset{c}{\downarrow} d. \tag{2}$$

From $ac \downarrow bd$ and $b \downarrow d$, Corollary II.2.10 yields

$$b \underset{ac}{\downarrow} d. \tag{3}$$

Since $b \rhd_a e_{ab}$, (1) and Lemma VI.3.12 imply:

$$b \rhd_{ac} e_{ab}. \tag{4}$$

Similarly $d \rhd_c e_{cd}$, (2), and Lemma VI.3.12 imply:

$$d \rhd_{ac} e_{cd}. \tag{5}$$

Now (3), (4), and (5) imply $e_{cd} \downarrow_{ac} e_{ab}$. Using $b \rhd_a e_{ab}$ again, we conclude from (1) that $ac \downarrow_a e_{ab}$. From transitivity of independence, we deduce $e_{cd} \downarrow_a e_{ab}$. Now $p_{cd} \dashv a$ implies $p_{cd} \perp p_{ab}$. This contradiction implies $b \not\downarrow d$.

Note that the roles of bd and ac are completely symmetric in this argument. Thus repeating the argument with reversed roles yields $a \not\downarrow c$.

Beyond the additional hypothesis there is a certain asymmetry between the arguments for i) and ii). In i) we use the fact that one side of the square is dependent in proving that the other is. In ii), with no hypothesis about either side, we prove by contradiction that a given side is dependent.

If we were able to conclude that an arbitrary 4-tuple $\langle a, b, c, d\rangle$ in normal position, with $p_{ab} \not\perp p_{cd}$ satisfied $a \not\downarrow c$ and $b \not\downarrow d$ then we could code a graph G inside a model M by requiring $dim(p_{ab}, M)$ to be small iff $\langle a, b\rangle \in G$. The necessity of the extra hypothesis in Proposition 3.12 ii) forces us to be more subtle and replace a single link by an infinite set of links. Moreover, in the case when $q_1 \neq q_2$ we must make a further reduction which follows from Proposition 3.12 i) when $q_1 = q_2$. From Proposition 3.10 we know that if four points are in normal position and the associated types are nonorthogonal then either the left side or the right side must be dependent. We show now that, in fact, this result holds uniformly.

3.13 Proposition. *Either*

i) *for every* $\langle a, b, c, d\rangle$ *in normal position with* $p_{ab} \not\perp p_{cd}$, $a \not\downarrow c$ *or*
ii) *for every* $\langle a, b, c, d\rangle$ *in normal position with* $p_{ab} \not\perp p_{cd}$, $b \not\downarrow d$.

Proof. If $q_1 = q_2$ this follows from Proposition 3.12 i). If $q_1 \perp q_2$ let $\langle a, b, c, d\rangle$ be in normal position with $p_{ab} \not\perp p_{cd}$ and suppose $a \not\downarrow c$. If i) fails we can, by applying an automorphism, assume it fails for $\langle a, b, c', d'\rangle$ for some choice of c', d'. That is, $p_{ab} \not\perp p_{c'd'}$ but $a \downarrow c'$ and $b \not\downarrow d'$. Choose $\langle \hat{c}, \hat{d}\rangle$ to realize $t(c', d'; ab)$ but with $\hat{c}\hat{d} \downarrow_{ab} cd$. Since $q_1 \perp q_2$ $\langle c, d, \hat{c}, \hat{d}\rangle$ is in normal position. Triviality of forking yields $c' \downarrow ab$ and so $\hat{c} \downarrow ab$. By transitivity of independence $\hat{c} \downarrow c$. Moreover $d \downarrow \hat{d}$. For, if not, we have $\hat{d} \not\downarrow d$ and (by choice of d' and $\hat{d}$) $\hat{d} \not\downarrow b$. But then $d \not\downarrow b$. Since $p_{ab} \not\perp p_{cd}$ and $p_{ab} \not\perp p_{\hat{c}\hat{d}}$, transitivity of nonorthogonality for regular types yields the required contradiction.

With the proposition in mind we can assume the following

3.14 Convention. Condition 3.13 ii) holds.

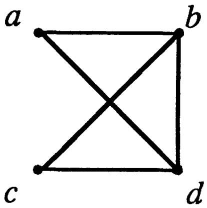

implies $\quad p_{ab} \perp p_{cd}$.

Fig. 12. Convention XVI.3.14.

We will not appeal to this convention until the proof of the last proposition in this section.

We next describe the family of graphs which we will be able to recover, then the properties of the model erected on each graph, and finally the recovery process.

3.15 Construction (The graph H). Let G be a connected symmetric graph on a set B. Add to B a single element $\hat{b}$ and let A be an infinite set with a distinguished element $\hat{a}$. Define a graph H on $A \cup B \cup \{\hat{b}\}$ as follows. Each element of B is connected to $\hat{a}$. Each element of A is connected to $\hat{b}$. If $b, b' \in B$ and $\langle b, b'\rangle \in G$ then each of b, b' is connected in H to every member of a countable subset $A_{b,b'}$ of A. The subsets of A associated with distinct pairs from B are disjoint. Rename $B \cup \{\hat{b}\}$ as B.

Thus the main feature relating G and H is that two elements, b_1, b_2, of B are connected in G if and only if there is an infinite subset of A such that each element of that subset is connected in H to both b_1 and b_2.

Take the set $A \subseteq q_1(\mathcal{M})$ and the set $B \subseteq q_2(\mathcal{M})$ such that $A \cup B$ is an independent set. Now, by Lemma II.2.26 extend $A \cup B$ to $D = A \cup B \cup E$ by choosing for each germane pair $\langle a, b\rangle$ from $A \times B$ an element e_{ab} such that $\langle a, b, e_{ab}\rangle$ is a germane triple and so that D is independent with respect to the partial order whose only relations are $a < e_{ab}$ and $b < e_{ab}$. If $q_1 = q_2$, choose $e_{ab} = e_{ba}$. Fix for $\langle a, b\rangle \in A \times B$, a copy p_{ab} of p in $S(e_{ab})$. If $q_1 = q_2$, let $p_{ab} = p_{ba}$.

3.16 Exercise. Show $\langle \hat{a}\hat{b}\rangle$ is a germane pair.

We will use the following abbreviation only in the next proposition. For $\langle a, b\rangle \in A \times B$, let $D_{ab} = \{a, b, e_{ab}\}$.

3.17 Proposition. *For $\langle a, b\rangle \in A \times B$, $p_{ab} \perp t(D - D_{ab}; D_{ab})$.*

Proof. By the choice of $<$, $(D - D_{ab}) \downarrow_{ab} D_{ab}$. But, $(D - D_{ab}) \downarrow ab$ so $(D - D_{ab}) \downarrow (D_{ab})$. Since $p_{ab} \dashv \emptyset$ we finish.

With this in hand we can define from H (and thus from G) a model M^G such that $M^G \approx M^{G'}$ implies $G \approx G'$.

3.18 Construction (The model M^G). Let $S = \{p_{ab} : \langle a, b\rangle \in H\}$. Form by Theorem XIV.3.5 a model M^G with $|M^G| = \lambda$ and containing D so that $dim(p, M^G) < \lambda$ if $p \in S$ and $dim(q, M^G) = \lambda$ if q is irrelevant to S.

We will recover H and thus G from M^G by recovering the equivalence classes under forking of the elements of $A \cup B$ and then seeing which of them are related by H.

3.19 Notation. We will refer to a pair $\langle a, b\rangle \in A \times B$ as a *standard pair.* For any $d \in Q(M^G)$ we denote the equivalence class of d (for the equivalence relation of forking restricted to $Q(M^G)$) by $[d]$. Since M^G is fixed, we write Q for $Q(M^G)$.

Note that if $c \in Q$, $[c] \cap (A \cup B)$ has at most one element. The remainder of the proof depends on the crucial observation that for any germane pair $\langle a, b\rangle$, if $dim(p_{ab}, M^G) < \lambda$ then for some standard pair $\langle a', b'\rangle$, $p_{ab} \not\perp p_{a'b'}$.

The first step is to recover the domain of the graph H from M^G.

3.20 Proposition (Recover domain).

i) *If $c \in Q$ then $[c] = [a]$ for some $a \in A$ iff $dim(p_{c\hat{b}}, M^G) < \lambda$.*
ii) *If $c \in Q$ then $[c] = [b]$ for some $b \in B$ iff $dim(p_{c\hat{a}}, M^G) < \lambda$.*

Proof. i) The construction guarantees that the condition is fulfilled whenever $[c] = [a]$ for some $a \in A$. If $dim(p_{c\hat{b}}, M^G) < \lambda$ then for some standard pair $\langle a, b\rangle$, $p_{c\hat{b}} \not\perp p_{ab}$. If $b \not\downarrow \hat{b}$ then $b = \hat{b}$ so by Proposition 3.7 i) we finish. Thus, $b \downarrow \hat{b}$. If $q_1 \neq q_2$, $q_1 \perp q_2$ and $\langle c, \hat{b}, a, b\rangle$ are in normal position so by Proposition 3.10, $a \not\downarrow c$. If $q_1 = q_2$ then $\langle b, a\rangle$ is a germane pair. Suppose for contradiction that $a \downarrow c$. Then $\langle c, \hat{b}, b, a\rangle$ is in normal position. By Proposition 3.12 i) $c \not\downarrow b$ and $\hat{b} \not\downarrow a$. But $a \in A$ and $\hat{b} \in B$ so this last condition is impossible and we finish.

ii) The proof of ii) is entirely analogous.

It remains to recover the graph H. We must introduce one last bit of notation.

3.21 Notation. For $b \in B$ and an infinite $X \subseteq A$, we write $X \mathcal{F} b$ if for all but finitely many elements x of X, $\langle x, b\rangle \in H$.

We now show that $\mathcal{F}$ can be recovered from M^G. It is clear that we can recover G from this $\mathcal{F}$.

3.22 Proposition. *For any $X \subseteq A$ and any $b \in B$, $X\mathcal{F}b$ if and only if*

$$\begin{array}{l} \textit{for some } \tilde{b} \in [b] \textit{ and all but finitely many } x \in X, \\ \textit{there is an } x' \in [x] \textit{ with } dim(p_{x'\tilde{b}}, M^G) < \lambda. \end{array} \tag{$*$}$$

Note that we are able to choose a fixed $\tilde{b} \in [b]$. This strengthens the natural version which would make $\tilde{b}$ depend on x.

Proof. Clearly $X\mathcal{F}b$ implies $(*)$. Suppose $(*)$ holds for X and b but $X\mathcal{F}b$ fails. Then there is an infinite subset $X_0 \subseteq X$ such that for each $x \in X_0$ there is an $a_x \in [x]$ with $dim(p_{a_x\tilde{b}}, M^G) < \lambda$ but $dim(p_{xb}, M^G) = \lambda$. The first condition implies that for each $x \in X_0$ there is an $\hat{x} \in A$ and a $b_x \in B$ such that $p_{a_x\tilde{b}} \not\perp p_{\hat{x}b_x}$. Since $A \cup B$ is independent, $A \cap B = \emptyset$, and forking is transitive on $Q(M^G)$, the tuple $\langle a_x, \tilde{b}, \hat{x}, b_x\rangle$ is in normal position for each $x \in X_0$. By Convention 3.14 for each x, $\tilde{b} \not\downarrow b_x$. Since each $b_x \in B$, this implies all the b_x are equal. Call the common value $\bar{b}$.

Thus we have a single pair $\langle \tilde{b}, \bar{b}\rangle$ and an infinite set of pairs $\langle a_x, \hat{x}\rangle$ such that for each x, $p_{a_x\tilde{b}} \not\perp p_{\hat{x}\bar{b}}$.

Let $X_0' = \{a_x : x \in X_0\}$ and $\hat{X}_0 = \{\hat{x} : x \in X_0\}$. Clearly X_0' is infinite. To see that $\hat{X}_0$ is infinite we show the map $x \mapsto \hat{x}$ is $1-1$. If not, for some $x \neq y \in X_0$ we have $p_{a_x\tilde{b}} \not\perp p_{\hat{x}\bar{b}}$ and $p_{a_y\tilde{b}} \not\perp p_{\hat{x}\bar{b}}$ so $p_{a_x\tilde{b}} \not\perp p_{a_y\tilde{b}}$ which implies $a_x \not\downarrow a_y$, contrary to the hypothesis that $x \downarrow y$. Now choose $X_1 \subseteq X_0' \cup \hat{X}_0$ such that X_1 is finite and $\tilde{b}\bar{b} \downarrow_{X_1} X_0' \cup \hat{X}_0$. Then choose $w \in X_0$ such that neither a_w nor $\hat{w}$ is in X_1. Since A is independent, this implies $a_w\hat{w} \downarrow \tilde{b}\bar{b}$. Since $\langle a_w, \tilde{b}, \hat{w}, \bar{b}\rangle$ is in normal position and $p_{a_w\tilde{b}} \not\perp p_{\hat{w}\bar{b}}$, this contradicts Proposition 3.12 ii) and completes the proof of the proposition.

Since there are 2^λ nonisomorphic graphs of the requisite sort and we have proved that $M_G \approx M_{G'}$ implies $G \approx G'$, we conclude Theorem 3.1.

3.23 Historical Notes. The proof presented in this section is based on that in [Harrington & Makkai 1985]. This proof differs significantly from Shelah's [Shelah 1982]. Here we use the forking technology to construct an interpretation of the theory of graphs into models of T. Shelah makes another application of the general method discussed in Section I.5. That is, he constructs generalized Ehrenfeucht-Mostowski models over the graphs and argues on general grounds that they are not isomorphic. This general argument requires a detailed combinatorial analysis in Chapter VIII of [Shelah 1978]. We avoid that analysis. On the other hand, Shelah applies his method again to show that theories with the 'omitting types order property' [Shelah 198?] have the maximal number of models in each uncountable power. There is no known extension of the method here to that case. Shelah's method will be expounded in [Hodges 198?]. Saffe [Saffe 1982] also proved this result by a method similar to that used here. This result leaves one lacuna in the general program. Let T be a countable superstable theory which is not ω-stable and suppose the continuum hypothesis fails.

Then for an uncountable cardinal, κ, below the continuum we have not proved that the dimensional order property implies that T has 2^κ models of power κ. This problem is resolved by Theorem IX.1.20 of [Shelah 1978].