Chapter V

Indiscernibles In Stable Theories

Sets of indiscernibles play a number of important roles in model theory. They are used to realize types (as in the proof of the existence of saturated models); they are also used to 'blow up' models (without realizing new types). Moreover the cardinalities of maximal sets of indiscernibles can be used as invariants in classifying models. We begin this chapter by expounding the basic properties of indiscernibles in a stable theory and explaining the distinction between sequences of indiscernibles and independent sets. Section 2 begins the rather lengthy process of using sets of indiscernibles as bases for models of stable theories (cf. the introduction to Section 2). In Section 3 we apply the notion of indiscernibility to show the equivalence between the notion of forking as introduced here and the original version of Shelah [Shelah 1978].

1. *Sets Of Indiscernibles*

If X is a set of algebraically independent elements in an algebraically closed field then every permutation of X is in fact an elementary map. This indiscernibility of the elements of X is closely related to their independence. We explain this connection in Lemma 1.8 and Theorem 1.23. In this section we study in detail indiscernible elements in a model of a stable theory. We define such notions as indiscernible sequences (of sequences) and indiscernible sets (of sequences). That is, we deal with families $E = \{\overline{e}_i : i \in I\}$ where each $\overline{e}_i$ is a finite sequence. Very little intuition is lost by thinking of each $\overline{e}_i$ as a single individual but the added generality is necessary.

1.1 Definition. The ordered set $X = \{\overline{x}_i : i \in l\}$ is a *sequence of (order) indiscernibles* if every order preserving map f mapping into a finite subset of the linear order l induces a partial elementary monomorphism of X by taking $\overline{x}_i$ to $\overline{x}_{f(i)}$.

The index set will be well-ordered unless we explicitly assert otherwise. We distinguish now between a sequence and a set of indiscernibles. While this is an important distinction in general model theory, we will see that in our situations the concepts coalesce.

1.2 Definition. The set $X = \{\overline{x}_i : i \in I\}$ is a *set of (pure) indiscernible sequences* if every partial permutation of X is an elementary monomorphism.

In Definition 1.1 (1.2) we refer to X as a *sequence* (*set*) of indiscernibles. Naturally, both of these definitions can be made relative to a set A.

The next theorem emphasizes two important properties of indiscernible sequences in a stable theory.

1.3 Theorem. *Let M be a model of the stable theory T.*

i) *Every infinite sequence of indiscernibles in M is a set of indiscernibles.*
ii) *For any formula $\phi(\overline{x};\overline{y})$, there is an integer $n(\phi)$ such that for any $\overline{b}$, if $E = \{\overline{e}_i : i \in \omega\}$ is an infinite set of indiscernibles over the empty set either $|\{i : \phi(\overline{e}_i;\overline{b})\}| < n(\phi)$ or $|\{i : \neg\phi(\overline{e}_i;\overline{b})\}| < n(\phi)$.*

Proof. In showing i), let S_n denote the symmetric group on n elements and denote the application of a permutation $s \in S_n$ to an element i by si. Fix $i_1 < \ldots < i_n \in \omega$, and consider an arbitrary formula $\phi(\overline{x}_1, \ldots, \overline{x}_n)$. Let T_n be the set of $s \in S_n$ such that $\models \phi(\overline{e}_{i_{s1}}, \ldots, \overline{e}_{i_{sn}})$. If T_n is $\emptyset$ or S_n we finish. If not, fix $t \in T_n$ and $s \in S_n - T_n$. Then s and t differ by a product of transpositions so there exist $u \in T_n$ and $v \in S_n - T_n$ such that for some $k < n, v = (k, k+1)u$. (This requires a little argument with permutations.) Let $\psi(\overline{x}_1, \ldots, \overline{x}_n) = \phi(\overline{x}_{v1}, \ldots, \overline{x}_{vn})$. Now we have

a) $\models \psi(\overline{e}_{i_1}, \ldots, \overline{e}_{i_n})$ and
b) $\models \neg\psi(\overline{e}_{i_1}, \ldots, \overline{e}_{i_{k-1}}, \overline{e}_{i_{k+1}}, \overline{e}_{i_k}, \ldots, \overline{e}_{i_n})$.

Note that a) and b) hold when $i_1, \ldots, i_n$ are replaced by any other properly ordered n-tuple from ω (by order indiscernibility). Now by compactness, we can replace ω by an index set R with order type of the real numbers such that a) and b) hold whenever $i_1, \ldots, i_n$ are replaced by a properly ordered n-tuple $r_1, \ldots, r_n$ from R. If r and r' are arbitrary distinct members of R, there exist rational numbers $\{q_i : 1 \leq i \leq n \wedge i \neq k\}$ such that $q_1 < \ldots < q_{k-1} < r < q_{k+1} < r' < q_{k+2} < \ldots < q_n$. Hence

c) $\models \psi(\overline{e}_{q_1}, \ldots, \overline{e}_{q_{k-1}}, \overline{e}_r, \overline{e}_{q_{k+1}}, \ldots, \overline{e}_{q_n})$ and
d) $\models \neg\psi(\overline{e}_{q_1}, \ldots, \overline{e}_{q_{k-1}}, \overline{e}_{r'}, \overline{e}_{q_{k+1}}, \ldots, \overline{e}_{q_n})$.

Thus distinct $\overline{e}_r$ for $r \in R$ realize distinct ψ-types over the countable set of $\overline{e}_q$ indexed by rationals. By Theorem III.1.6, T is not stable.

For ii), suppose first there is an infinite and coinfinite subset F of E such that $\phi(\overline{e}, \overline{b})$ holds iff $\overline{e}$ is in F. For each infinite and coinfinite subset W of E, there is an elementary monomorphism f_w with $f(F) = W$. Since the monster model is so homogeneous, we can extend f to an automorphism of the monster model. Now, $\phi(\overline{a}_i, f(\overline{b}))$ holds iff $\overline{a}_i$ is in W. Thus there are continuum many ϕ-types over E, contrary to the stability of T.

This argument depended on the choice of F and b. However, if arbitrarily large finite F could be found for various choices of $\overline{b}$, by compactness we could obtain a single F and $\overline{b}$ to which the above proof applies.

Note that Theorem 1.3 ii), but not necessarily i), holds for any theory which does not have the independence property (cf. Section III.4.43).

1.4 Exercise. Find an example showing the necessity of the assumption that the sequence is infinite in Theorem 1.3 i).

By Theorem 1.3 i) in stable theories all infinite sequences of indiscernibles are sets of pure indiscernibles. In fact the converse is true, but we won't prove it here (see Chapter 2 in [Shelah 1978]). Using the second part of this theorem we can attach to each set of indiscernible n-tuples a complete n-type as follows.

1.5 Definition. Let $X = \{\overline{a}_i : i \in I\}$ be an infinite set of indiscernibles over the empty set. The *average type* of X over A, denoted $\mathrm{Av}(X; A)$, is $\{\phi(\overline{x}; \overline{a}) : \phi(\overline{x}; \overline{a})$ in $F(A)$ such that for all but finitely many i, $\models \phi(\overline{a}_i; \overline{a})\}$.

Thus, if X is an indiscernible set of n-tuples, $\mathrm{Av}(X, A) \in S^n(A)$. Note that if X is indiscernible over any set B it is certainly indiscernible over the empty set; X is usually not indiscernible over A when we apply this definition.

1.6 Exercise. Show that if $\overline{e}$ realizes $\mathrm{Av}(E; A)$ where E is an infinite set of indiscernibles over $B \subseteq A$ and $B \cup E \subseteq A$ then $E \cup \overline{e}$ is a set of indiscernibles over B.

The identification of indiscernibles with independent elements in vector spaces, algebraically closed fields, and free algebras is not a fluke. These examples are somewhat misleading, however. Here are some more revealing ones.

1.7 Examples. First we show a sequence may be indiscernible without being independent.

i) Let T be the theory of an equivalence relation with infinitely many infinite classes. If X is an infinite set of points in the same equivalence class then X is a set of indiscernibles but X is not an independent set.

Here is an independent set which is not a set of indiscernibles.

ii) In the theory of infinitely many refining equivalence relations with finite splitting (REF_ω) let $\langle a_i : i < \omega \rangle$ be a sequence such that $E_i(a_k, a_j)$ iff $k, j \geq i$. Then for each i, $t(a_i; A_i)$ does not fork over $\emptyset$ but the sequence $\langle a_i : i < \omega \rangle$ is not indiscernible over the empty set. Note also that $t(a_i; A_i) \subseteq t(a_j; A_i)$ if $i < j$.

These examples show that independence does not quite guarantee indiscernibility; however, only one ingredient is missing. Recall Definition II.2.19 declares that an independent sequence is strongly independent over A if all elements of the sequence realize the same stationary type over A.

1.8 Lemma. *If E is an infinite strongly independent set over A then E is a set of indiscernibles over A.*

Proof. Fix a well ordering $\langle \overline{a}_\alpha : \alpha < \kappa \rangle$ of E. We prove by induction on $\beta < \kappa$ that E_β is an indiscernible sequence over A; by Theorem 1.3 this suffices. The result is clear for limit ordinals, so suppose $\beta = \alpha + 1$ and E_α is a sequence of indiscernibles over A. Let $\overline{a}_{i_0}, \ldots, \overline{a}_{i_n}$ and $\overline{a}_{j_0}, \ldots, \overline{a}_{j_n}$ be in E_β with $i_0 < i_1 < \ldots < i_n$, $j_0 < j_1 < \ldots < j_n$, and without loss of generality $i_n < j_n$. We need only show

$$t(\overline{a}_\alpha; \{\overline{a}_{i_0}, \ldots, \overline{a}_{i_n}\} \cup A) = t(\overline{a}_\alpha; \{\overline{a}_{j_0}, \ldots, \overline{a}_{j_n}\} \cup A).$$

By induction, there exists a map f which fixes A and sends $\overline{a}_{i_l}$ to $\overline{a}_{j_l}$ for $l \leq n$. Both $t(\overline{a}_\alpha; \{\overline{a}_{j_0}, \ldots, \overline{a}_{j_n}\} \cup A)$ and $f(t(\overline{a}_\alpha; \{\overline{a}_{i_0}, \ldots, \overline{a}_{i_n}\} \cup A))$ are nonforking extensions of the same stationary type so they are equal. Applying f^{-1} we have the result.

Theorem 1.23 below is almost a converse to this result. We can weaken the hypothesis that E is strongly independent over A slightly as follows. The second exercise shows that this provides a real weakening of the hypothesis.

1.9 Exercise. Show that if the infinite set E is independent over A and all $\overline{e} \in E$ realize the same strong type over A then E is a set of indiscernibles over A. (Hint: First apply Lemma 1.8 to show $E - \{\overline{e}\}$ is a set of indiscernibles over $A \cup \{\overline{e}\}$ for each $\overline{e} \in E$. Then deduce that E is a set of indiscernibles over A.)

1.10 Exercise. Let M be a model of CEF_ω and suppose $\{c_i : i < \omega\}$ is an infinite set of elements realizing the same strong type over $\emptyset$. Show using the previous exercise (or by inspection) that the c_i are indiscernible. Show also that no type over the empty set is stationary.

The argument for Lemma 1.8 does not depend on the particular properties of stable theories but relies only on the definitions of stationary type and strongly independent sequence. We will use stability theory to construct a strongly independent sequence. We first describe a somewhat more syntactic way of finding an indiscernible sequence and indicate how to find indiscernibles in models of arbitrary theories. We devote the remainder of the section to elucidating those properties of sets of indiscernibles which depend on stability.

The following exercises are not necessary for our development but they put our uses of the hypothesis of stability into context. For this set of exercises we drop the standing assumption that T is stable.

1.11 Exercise. Show that if $\langle p_i : i < \alpha \rangle$ is an increasing sequence of types, each $p_i \in S(A \cup E_i)$, with $A \subseteq \operatorname{dom} p_0$, no p_i splits (Definition III.2.9) over A, and e_i realizes p_i then $\langle e_i : i < \alpha \rangle$ is a sequence of indiscernibles.

Using this observation and the fact that coheirs over models always exist solve the following exercises for an arbitrary theory T.

1.12 Exercise. If M is a model of T, find in the monster model an infinite set of indiscernibles over M.

1.13 Exercise. If $M \subseteq N$ and N is $|M|^+$-saturated find in N an infinite set of indiscernibles over M.

1.14 Exercise. If $A \subseteq N$ and N is $(|A| + |T|)^+$-saturated find in N an infinite set of indiscernibles over A.

The remaining results in this section depend, at least indirectly, on the hypothesis that types over models are stationary. The essential point is that indiscernible elements in stable theories are even more similar than is evident from the definition of indiscernibility. We first show that if E is a set of indiscernibles over A (i.e. the elements of E cannot be distinguished by formulas which are over A) then in fact the elements of E cannot be distinguished by formulas which are almost over A.

Note that in the following lemma E is a sequence of sequences. Thus, each variable $\overline{x}$ and each constant $\overline{e}$ refers to a sequence of sequences, e.g. $\overline{x} = \langle \overline{y}_0, \ldots, \overline{y}_m \rangle$, where $\lg(\overline{y}_i)$ is the common length of the sequences in E.

1.15 Lemma. *If E is an infinite set of sequences which is indiscernible over A, and $\phi(\overline{x}_1, \ldots, \overline{x}_n; \overline{c})$ is almost over A then every sequence $\langle \overline{e}_1, \ldots \overline{e}_n \rangle$ from E yields the same truth value for $\phi(\overline{e}_1, \ldots, \overline{e}_n; \overline{c})$.*

Proof. If the lemma fails, we can choose disjoint sequences (of sequences) $\overline{h}$ and $\overline{g}$ from E such that $stp(\overline{h}; A) \neq stp(\overline{g}; A)$. (If $\overline{h}$ and $\overline{g}$ are not disjoint replace one of them by a sequence disjoint from both of the original two.) There is a member R of $FE^m(A)$ such that $\neg R(\overline{g}, \overline{h})$. Thus any pair of disjoint sequences from E each having the same length as $\overline{g}$ are inequivalent under R so R has infinitely many equivalence classes, contrary to the choice of R in $FE^m(A)$.

1.16 Exercise. Show that if E is an infinite set of indiscernibles over A then E is an infinite set of indiscernibles over $\mathrm{cl}(A)$. Show that it is essential to assume that E is infinite.

1.17 Exercise. Suppose $\theta(\overline{x})$ is almost over $B \subseteq A$ and $E \subseteq A$ is an infinite set of indiscernibles over B such that all but finitely many $\overline{e} \in E$ satisfy θ. Show that if $\overline{e}'$ realizes $\mathrm{Av}(E, A)$ then $\models \theta(\overline{e}')$.

Now we apply Lemma 1.15 to show that if E is a set of indiscernibles over A then E remains indiscernible over supersets of A which are independent from E over A. The next theorem does all the work. This is one of the most important facts about the nonforking relation.

1.18 Theorem. *If E is an infinite set of indiscernibles over A and if $t(\overline{b}; E \cup A)$ does not fork over A then E is a set of indiscernibles over $A \cup \overline{b}$.*

Proof. It suffices (since E is indiscernible over A) to show that the type q of an arbitrary n-tuple of elements from E has a unique extension $q' \in S(A \cup \overline{b})$ which is realized in E. But if $r \in S(A \cup \overline{b})$ is the type of an n-tuple from

E, by the symmetry lemma r does not fork over A. Thus, if E is not a set of indiscernibles over $A \cup \bar{b}$, we have two extensions, r and r' of q which are realized in E and do not fork over A. By a form (Theorem IV.1.23) of the Finite Equivalence Relation Theorem, this implies that r and r' imply contradictory formulas which are almost over A. This contradicts Lemma 1.15.

The proof of Lemma 1.15 does not rely on stability. However, to make effective use of Lemma 1.15 and Theorem 1.18 we need to have $\kappa(T) < \infty$. Some simple cardinal computations and the definition of $\kappa(T)$ yield the following consequences.

1.19 Corollary. *Let E be indiscernible over A.*

i) *For any $\bar{b}$, there is an $E_0 \subseteq E$ with $|E_0| < \bar{\kappa}(T)$ such that $E - E_0$ is indiscernible over $A \cup E_0 \cup \{\bar{b}\}$.*
ii) *For any B there is a subset $E_0 \subseteq E$ such that $E - E_0$ is indiscernible over $E_0 \cup A \cup B$ with*
   a) *$|E_0| \leq \kappa(T) + |B|$ and*
   b) *if $|B| < \kappa(T)$ then $|E_0| < \kappa(T)$.*

Proof. For i) use the definition of $\bar{\kappa}(T)$; for ii) we iterate the construction. For each finite sequence $\bar{b}$ from B choose $E(\bar{b})$ contained in E such that $|E(b)| < \kappa(T)$ and $E - E(\bar{b})$ is indiscernible over $A \cup \{\bar{b}\}$. Then, let E_0 be $\bigcup\{E(b) : b \text{ a finite sequence from } B\}$. It is easy to check E_0 has the required cardinality.

If we are willing to replace the $\bar{\kappa}(T)$ in i) by $|T|^+$, an argument both more naive and more general will suffice.

1.20 Exercise. Use Theorem 1.3 to show that if T does not have the independence property (cf. Paragraph III.4.42) and E is a set of indiscernibles over A then for any $\bar{b}$ there exists $E_0 \subseteq E$ with $|E_0| \leq |T|$ such that $E - E_0$ is indiscernible over $A \cup E_0 \cup \bar{b}$.

Our next aim is to prove a suitable converse to Lemma 1.8. Recall the definition from II.2.14 of a coherent sequence.

1.21 Lemma. *Let $E = \{\bar{e}_i : i < \beta\}$ be a coherent sequence over A.*

i) *If $\beta \geq \omega$ then all $\bar{e}_m$ for $m \geq \omega$ realize the same strong type over A.*
ii) *If $|E| \geq \bar{\kappa}(T)$ then E contains a subset E_0 with $|E_0| < \kappa(T)$ such that $E - E_0$ is independent over $A \cup E_0$.*

Proof. i) For each finite equivalence relation R over A, there is an n (one more than the number of classes of R) such that for all $\alpha \geq n$, $\bar{e}_\alpha$ is R-equivalent to some (the same) member of E_n. Thus all $\bar{e}_\alpha$ for $\alpha \geq \omega$ satisfy the same strong type over A.

ii) There exists $E_0 \subseteq E$ with $|E_0| < \bar{\kappa}(T)$ such that $\mathrm{Av}(E; E \cup A)$ does not fork over $E_0 \cup A$. Since E is coherent, for any $\bar{e} \in E - E_0$, $t(\bar{e}_\alpha; E_\alpha \cup A)$ is a restriction of $\mathrm{Av}(E; E \cup A)$ so we finish.

The property described in Lemma 1.21 i) depends on the order type rather than the cardinality of the sequence E. Note that Lemma 1.21 i) does not depend on the cardinality of L. It holds for a sequence E, even if $|L|$ is much greater than $|E|$.

1.22 Exercise. Show that even for countable T, there can be coherent sequences over a set A of arbitrary finite length $\langle e_i : i < n \rangle$ such that not all the e_i realize the same strong type over A.

Since any sequence of indiscernibles is a coherent sequence we could apply Lemma 1.21 directly to any indiscernible set E to obtain a proper subset E_0 of E over which $E - E_0$ is strongly independent and with $|E_0| < \kappa(T) + \aleph_1$. However, we can use Lemma 1.15 to improve the bound on $|E_0|$. In particular, if T is superstable we choose E_0 in the following theorem to be finite.

1.23 Theorem. *If E is a set of indiscernibles over A and $|E| \geq \overline{\kappa}(T)$ then there is a set $E_0 \subseteq E$ with $|E_0| < \overline{\kappa}(T)$ such that $E - E_0$ is strongly independent over $E_0 \cup A$.*

Proof. Since E is a coherent sequence, we can choose by Lemma 1.21 ii) a set E_0' with $|E_0'| < \overline{\kappa}(T)$ such that $E - E_0'$ is independent over $A \cup E_0'$. Let $E_0 = E_0' \cup \overline{e}$ for any $\overline{e}$ in $E - E_0'$. By Lemma 1.15 all elements of $E - E_0'$ realize the same strong type over $A \cup E_0'$ and this strong type is specified by the choice of $\overline{e}$.

We have discussed the notions: a) a strongly independent sequence b) an independent sequence and c) a sequence of indiscernibles. We have shown that a) implies both b) and c). Neither converse holds. But, Theorem 1.23 asserts that each set satisfying c) 'eventually' satisfies a).

From Lemma 1.15 we see that if E is a set of indiscernibles over A then all elements of E realize the same strong type over A. Moreover, for any k, any k-tuple of sequences from E realizes the same strong type over A. For applications, we need a somewhat stronger result. Namely, we would like the type over A to actually specify which strong type is realized by the sequence of indiscernibles. To arrange this for 1-tuples from E is easy; we just add one element of E to A and work over $A \cup \overline{e}$. It would seem that to deal with k-tuples from E we would have to add k members of E. This would be unfortunate if E had only $\aleph_0$ members. However, here we are able to use the strength of independence instead of mere indiscernibility. The case of the following theorem with $\lg(\overline{h}) = 2$ is Exercise II.2.1. The argument here requires one crucial observation about automorphisms and the equality of types: If α is an automorphism fixing B and taking $\overline{a}$ to $\overline{a}'$ and $t(\alpha\overline{b}; \overline{a}' \cup B) = t(\overline{b}'; \overline{a}' \cup B)$ then $t(\overline{b}^\frown\overline{a}; B) = t(\overline{b}'^\frown\overline{a}'; B)$.

1.24 Theorem. *If E is strongly independent over A then for any sequence $\overline{h}$ of sequences from E (i.e. $\overline{h} = (\overline{e}_{i_0}, \ldots, \overline{e}_{i_n})$), $t(\overline{h}; A)$ is stationary.*

Proof. The proof is by induction on the length of $\overline{h}$. If $\lg(\overline{h}) = 1$, this is the definition of strong independence. Now suppose $\overline{h} = \overline{g}^\frown\overline{e}$ and $t(\overline{g}; A)$

is stationary. Suppose that there exist $B \supseteq A$, $\overline{g}'$, $\overline{g}''$, $\overline{e}'$, $\overline{e}''$ such that $t(\overline{g}'^\frown\overline{e}';A) = t(\overline{g}''^\frown\overline{e}'';A) = t(\overline{g}^\frown\overline{e};A)$ and $t(\overline{g}'^\frown\overline{e}';B)$ does not fork over A and $t(\overline{g}''^\frown\overline{e}'';B)$ does not fork over A. Since $t(\overline{g}'^\frown\overline{e}';A) = t(\overline{g}^\frown\overline{e};A)$, the independence of E implies $t(\overline{g}';A \cup \overline{e}')$ does not fork over A. But then, by Corollary II.2.10, $t(\overline{g}';B \cup \overline{e}')$ does not fork over B. Since $t(\overline{e}';B)$ does not fork over A and $t(\overline{e}'';B)$ does not fork over A and their common restriction to A is stationary, $t(\overline{e}';B) = t(\overline{e}'';B)$. Hence, there is an automorphism α which fixes B and maps $\overline{e}''$ to $\overline{e}'$. We can apply to $\alpha\overline{g}''$ the same argument we applied to $\overline{g}'$ and conclude $t(\alpha\overline{g}'';B \cup \overline{e}')$ does not fork over A. Thus by the stationarity of $t(\overline{g};A)$, $t(\overline{g}';B \cup \overline{e}') = t(\alpha\overline{g}'';B \cup \overline{e}')$. Hence we can conclude $t(\overline{g}'^\frown\overline{e}';B) = t(\overline{g}''^\frown\overline{e}'';B)$ as required.

1.25 Exercise. Produce an example where $t(\overline{a};A \cup \overline{b})$ is stationary and $t(\overline{b};A \cup \overline{a})$ is not. Refine the example to make $\overline{a} \downarrow_A A \cup \overline{b}$.

For simplicity of notation, we state the following theorem for a set of indiscernibles over the empty set.

1.26 Corollary. *If T is a small superstable theory and I is a countable set of indiscernibles over $\emptyset$, then $|S(I)| = \aleph_0$.*

Proof. Suppose $S(I)$ is uncountable. For every $p \in S(I)$ there is a finite $I_p \subseteq I$ such that p does not fork over I_p. Thus, there is an $I_0 \subseteq I$ with $2^{\aleph_0}$ types $p \in S(I)$ which do not fork over I_0. Without loss of generality, $I - I_0$ is strongly independent over I_0. Let, then, $\overline{a}$ and $\overline{b}$ realize the same type over I_0 but distinct types over I with $a \downarrow_{I_0} I$ and $b \downarrow_{I_0} I$. Thus, for some finite sequence $\overline{e} \in I$ we have $t(\overline{e};I_0 \cup \overline{a}) \neq t(\alpha(\overline{e});I_0 \cup \overline{a})$ where α is an automorphism of $\mathcal{M}$ which fixes I_0 and maps $\overline{b}$ to $\overline{a}$. But both $\overline{e} \downarrow_{I_0} \overline{a}$ and $\alpha(\overline{e}) \downarrow_{I_0} \overline{a}$ so this contradicts Theorem 1.24, which implies that $t(\overline{e};I_0)$ is stationary.

1.27 Exercise. Suppose $\langle p_i : i < \alpha\rangle$ is an increasing sequence of types with the ordinal $\alpha > \omega$ and $\mathrm{cf}(|\alpha|) > \kappa(T)$. If $E = \{e_i : i < \alpha\}$ where e_i realizes p_i then E contains an infinite indiscernible subsequence.

Thus far in this section we have dealt with indiscernibles without generally guaranteeing their existence or their location. In Exercises 1.13 and 1.14 we showed the existence of indiscernibles somewhere in the monster model; now, we employ a variant on an argument due originally to Morley to construct indiscernibles within a given set. The specific argument here was suggested to me by Lascar.

1.28 Theorem. *Let T be λ-stable, $\lambda \geq |T| + \kappa(T)$, and A a subset of B with $|B| > \lambda \geq |A|$. Then B contains a subset E with $|E| > \lambda$ which is strongly independent over some M containing A and thus indiscernible over A.*

Proof. (Fig. 1). Choose by induction sets $E_i \subseteq B$ such that E_{i+1} is maximal with respect to being independent over $A \cup E_i$ (and taking unions at limits). If for each i, $|E_i| \leq \lambda$ this construction can be continued at least λ times.

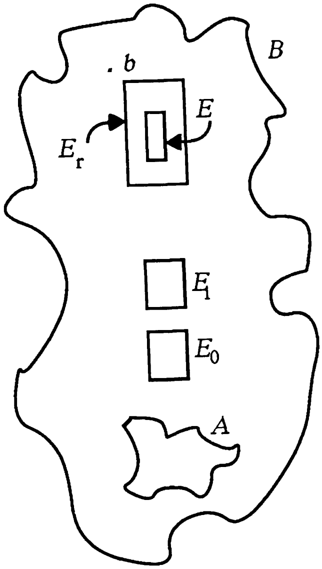

Fig. 1. Theorem 1.28. Constructing Indiscernibles

Choose any $b \in B - \bigcup_{i<\lambda} E_i$. Then, for each i, by the definition of E_i, $t(b; A \cup \bigcup_{j \le i} E_j)$ forks over $\bigcup_{j<i} E_j \cup A$. Since $\lambda \ge \kappa(T)$, this is impossible. Thus for some $k < \lambda$, $|E_k| > \lambda$. Choose M containing A such that $|M| = \lambda$ and $t(M; B)$ does not fork over $\bigcup_{j<k} E_j \cup A$. Since T is stable in λ there is a subset E of E_k with $|E| > \lambda$ such that all elements of E realize the same, necessarily stationary, type over M. By the symmetry axiom and Corollary II.2.10, E is independent over M, whence by Theorem 1.3 indiscernible over M and a fortiori indiscernible over A.

1.29 Exercise. Derive Morley's theorem which states that if T is a countable ω-stable theory and A is an uncountable set then for any countable set $B \subseteq A$, A contains an uncountable set of elements indiscernible over B.

These ideas can be applied along with König's lemma to give the following results of [Zilber 1980b]. For the following three exercises, assume T is $\aleph_0$-categorical and superstable and so ω-stable. Let X be a finite subset of a model of T.

1.30 Exercise. Any infinite subset of a model of T contains an infinite set of indiscernibles over X. (Hint: use König's lemma and Lemma 1.21.)

Now using Theorem 1.18 deduce the following.

1.31 Exercise. If $X \subseteq A$, Y is infinite and $t(Y; A)$ does not fork over X then Y contains an infinite set of indiscernibles over A.

1.32 Exercise. If X is a subset of A, M is a minimal prime model over A (cf. Chapter IX), and Y is a subset of M which satisfies '$t(Y; A)$ does not fork over X', then Y is finite.

1.33 Historical Notes. The vital notion of a set of indiscernibles was introduced by Ehrenfeucht and Mostowski [Ehrenfeucht & Mostowski 1956] who used it to prove that a first order theory has models of arbitrary cardinality κ whose automorphism groups have cardinality 2^κ. In [Morley 1965], Morley discovered the important distinction between order and pure indiscernibles and introduced the inductive construction of sequences of indiscernibles. Shelah's recognition that this construction could be carried out with a more general technique than the minimization of Morley rank was one of the impulses towards the development of the nonforking notion. Shelah introduced the notion of an average type in [Shelah 1978]. Precursors of the important fact (Theorem 1.18) that nonforking extensions preserve indiscernibility are in [Harnik & Ressayre 1971] and [Shelah 1972].

2. *Comparing Sets of Indiscernibles*

We will eventually assign as invariants of a model the cardinalities of various sets of indiscernibles. While the structure of a vector space is determined by one magnitude or dimension, models of other theories will depend on more than one dimension. Consider first the theory of an infinite set and second the theory of an equivalence relation with two infinite classes. A model of the first is specified by one magnitude, its cardinality; a model of the second requires two magnitudes to be specified, the cardinality of each equivalence class. Once we agree to count more than one set of indiscernibles, we face a problem of redundant information. Of which sets of indiscernibles do we need to know the cardinality? For example, consider the theory of an equivalence relation with two infinite classes, each of which is a model of the theory of the integers under successor. Clearly, to determine a model we need to count the number of components in each class. In a particular class, a set of indiscernibles containing one element from each component and a second set consisting of the successors of the elements in the first set give the same information. (Fig. 2).

Fig. 2. The x's and o's give the same information.

Our aim, which can be achieved only in some cases, is to define an equivalence relation on sets of indiscernibles in a model, assign a dimension to each class, and have these dimensions determine the model up to isomorphism. For this procedure to be useful we need to check two things. First, that the obvious way of assigning such a dimension as the cardinality of a maximal set of indiscernibles in the equivalence class is well-defined. Second, that in fact the cardinalities of these maximal sets of indiscernibles determine the model up to isomorphism.

There are several difficulties in this program. The first condition is easily met for equivalence classes whose dimension is greater than $|T|$ and even without too much difficulty for those with cardinality above $\kappa(T)$. It is much more difficult for small (e.g. finite) dimensions. Unfortunately , it is just the finite dimensions which are crucial in investigating countable models and this is one reason for the difficulty of settling Vaught's conjecture. It is to meet this difficulty that the apparatus of regular types is developed in Chapter XII.

There is however a second difficulty with redundancy. It turns out that the equivalence relation suggested by the example beginning this section is not sufficiently coarse. If we were to expand that example by adding a 1-1 function between the two equivalence classes then we would need only one cardinal to specify a model instead of two. But, the indiscernibles in distinct classes are not equivalent. We will need to describe a second equivalence relation to allow for this second level of redundancy. This equivalence relation is defined in Chapter XI.

2.1 Definition. Let X_1 and X_2 be infinite sets of indiscernibles over A. We write $X_1 \sim_A X_2$ and say X_1 and X_2 are *equivalent* if there is an infinite set, Y, of indiscernibles over A such that $X_1 \cup Y$ and $X_2 \cup Y$ are sets of indiscernibles over A. We omit the A if it is the empty set.

Note that E_0 and E_1 can be equivalent without $E_0 \cup E_1$ being a set of indiscernibles. Just choose two distinct sets of representatives for the equivalence classes in the example considered in the paragraph before the definition. For another example, in the theory of the integers under successor let E_0 and E_1 each select a different element from each of infinitely many components of a model.

If $X_1 \sim_A X_2$ and $|X_1| = |X_2|$ then $t(X_1; A) = t(X_2; A)$ so there is an automorphism fixing A taking X_1 to X_2. Clearly $\sim_A$ is symmetric and reflexive; but to establish transitivity we prove the following characterization.

2.2 Theorem. *If X_1 and X_2 are infinite sets of indiscernibles over the empty set, the following are equivalent.*

i) $X_1 \sim X_2$.
ii) *For every A,* $\mathrm{Av}(X_1; A) = \mathrm{Av}(X_2; A)$.
iii) *For some model M containing $X_1 \cup X_2$,* $\mathrm{Av}(X_1; M) = \mathrm{Av}(X_2; M)$.

Proof. Suppose i) holds and $X_1 \cup Y$ and $X_2 \cup Y$ are both sets of indiscernibles over the empty set. Then for any A, since X_1, X_2, and Y are infinite, $\mathrm{Av}(X_1; A) = \mathrm{Av}(X_1 \cup Y; A) = \mathrm{Av}(Y; A) = \mathrm{Av}(X_2 \cup Y; A) = \mathrm{Av}(X_2; A)$. Thus ii) holds.

Now we show that ii) implies i). For this, first define by induction a set $C = \{\overline{c}_i : i \in \omega\}$ so that if D_n denotes $A \cup X_1 \cup X_2 \cup C_n$, $\overline{c}_n$ realizes $\mathrm{Av}(X_1; D_n)$. Observe that if the construction were repeated replacing $\mathrm{Av}(X_1; D_n)$ with $\mathrm{Av}(X_2; D_n)$ the resulting set C' would be isomorphic to C over $X_1 \cup X_2 \cup A$. Now, applying the fact that if F is a set of indiscernibles over an arbitrary set A and $\overline{c}$ realizes $\mathrm{Av}(F; A \cup F)$ then $F \cup \{\overline{c}\}$ is a set of indiscernibles over A, it is easy to show by induction that C is the required set of indiscernibles for the theorem.

Now we show ii) and iii) are equivalent. It is obvious that ii) implies iii). If for some A, $\mathrm{Av}(X_1; A) \neq \mathrm{Av}(X_2; A)$ then there is a formula $\phi(\overline{x}; \overline{a})$ such that $\phi(\overline{x}; \overline{a})$ is in $\mathrm{Av}(X_1, A)$ but $\neg\phi(\overline{x}; \overline{a})$ is in $\mathrm{Av}(X_2, A)$. Let $n = n(\phi)$ from Theorem 1.3 ii). Now choose $\overline{b}_1, \ldots, \overline{b}_n$ from X_1 and $\overline{c}_1, \ldots, \overline{c}_n$ from X_2 such that $\models \phi(\overline{b}_i; \overline{a}) \wedge \neg\phi(\overline{c}_i; \overline{a})$ for all $i < n+1$. Then since M is a model of T which contains both X_1 and X_2, for some a' in M,

$$\bigwedge_{i<n+1} \phi(\overline{b}_i; \overline{a}') \wedge \neg\phi(\overline{c}_i; \overline{a}')$$

is true. By Theorem 1.3, $\mathrm{Av}(X_1; M) \neq \mathrm{Av}(X_2; M)$.

Since ii) is obviously transitive, so is i) as required.

We want now to define the dimension of an equivalence class in this relation.

2.3 Definition. i) Let $A \subseteq B$, $E \subseteq B$ and suppose E is a set of indiscernibles over A. Then $\dim(E, A, B)$, the dimension of E in B over A, is the minimal cardinality of a maximal set of indiscernibles $F \subseteq B$ such that $E \sim_A F$.

We can extend this notion to types as follows.

ii) Let $p \in S(A)$ be a stationary type. If $A \subseteq B$, then $\dim(p, B)$ is the cardinality of a maximal independent coherent sequence over A in B of realizations of p.

The next lemma shows that for sufficiently large values of the dimension, it can be calculated as the size of *any* maximal independent set. This is analogous to the situation in linear algebra where a simple cardinality argument shows that the dimension of a vector space is well defined if it is infinite but a much more complicated argument is necessary for the case of finite dimension. In our situation, we will be unable to handle the case of finite dimension until Chapter XII when we restrict the class of sets of indiscernibles that we discuss.

2.4 Theorem. *Let E be a maximal set of indiscernibles over A with $A \cup E$ contained in B. If $\dim(E, A, B) \geq \kappa(T)$ then for any maximal indiscernible set $F \subseteq B$ with $F \sim_A E$, $|E| = |F|$.*

Proof. Without loss of generality, suppose for contradiction that $|F| > |E|$. Then by Corollary 1.19 we can find a subset F' of F with $|F'| > |E|$ such that F' is a set of indiscernibles over E. We will show $F' \cup E$ is a set of indiscernibles, contradicting the maximality of E. Since E and F are equivalent, there is an infinite set E_0 such that $F \cup E_0$ and $E \cup E_0$ are both indiscernible sets. Thus $F' \cup E_0$ is a set of indiscernibles and, possibly first extending E_0 and applying Corollary 1.19 again, we can assume $F' \cup E_0$ is indiscernible over E. But now it is easy to see that $F' \cup E$ is indiscernible over A.

Let E be a set of indiscernibles over A. The remainder of this section is devoted to the search for a syntactic description of E which preserves equivalence over A.

One description of E is $t(E; A)$. This description does not contain enough information to specify E up to equivalence. To see this, consider a theory T with two unary predicates R and S which partition the universe of each model. Further, let f be a function from R onto S such that each element has infinitely many preimages. Now consider sets X and Y of elements realizing R such that all elements of X map to a single a satisfying S; all elements of Y map to a single element b in S; $a \neq b$. Then each of X and Y is a set of indiscernibles and $t(X; \emptyset) = t(Y; \emptyset)$ but X is not equivalent to Y. For, suppose Z were an infinite set such that both $X \cup Z$ and $Y \cup Z$ were indiscernible. For any $z \in Z$, we must have $f(z) = a$ from the viewpoint of X and $f(z) = b$ from the viewpoint of Y.

We will describe E in terms of the type of a single element in the sequence E. To describe the cases in which this is possible, we require some futher definitions. These notions extend, but do not conflict with, the definition of a type being based or strongly based on a set (Definition IV.2.4).

2.5 Definition. i) The infinite set of indiscernibles E is *based* on A if for every B which contains A, $\mathrm{Av}(E, B)$ does not fork over A.

ii) If, in addition $\mathrm{Av}(E, A)$ is stationary then E is *strongly based* on A. If $q = \mathrm{Av}(E, A)$, we may write E is based on q.

iii) If $p \in S(A)$ and for some $\bar{c}$ realizing p, E is strongly based on q, we also say E is based on p.

Of course, either of the conditions in Definition 2.5 can be checked by considering $\mathrm{Av}(E, \mathcal{M})$ rather than all B containing A.

2.6 Lemma. *If E is an infinite indiscernible set then there is a definition d of $\mathrm{Av}(E, E)$ with parameters from E such that d defines $\mathrm{Av}(E, \mathcal{M})$.*

Proof. For each formula $\phi(\bar{x}; \bar{y})$ there is by Theorem 1.3 an $n(\phi)$ such that for any $\bar{c} \in \mathcal{M}$, $\phi(\bar{x}; \bar{c}) \in \mathrm{Av}(E, \mathcal{M})$ iff $\models \phi(\bar{e}_i; \bar{c})$ for at least $n = n(\phi)$ sequences $\bar{e}_i$ from E. Thus we can take as $d\phi$ the formula

$$\bigvee_{\substack{I \subseteq E_{n+1} \\ |I| = n}} \bigwedge_{i \in I} \phi(\bar{e}_i; \bar{y}).$$

We have, in fact, that $\mathrm{Av}(E, E)$ is defined by a positive Boolean combination of instances of ϕ. This result also holds for any type over a model (cf. [Baldwin & Shelah 1985] IV.2.15). This observation is also closely connected to the normalization lemma (cf. [Harnik & Harrington 1984] [Pillay

198?]). This similarity in role between a set of indiscernibles and a model appears again in Section X.4.

2.7 Theorem. *Let E be an infinite set which is indiscernible over A. Then E is strongly based on E.*

Proof. By Lemma 2.6, $\mathrm{Av}(E; \mathcal{M})$ is definable over E so $\mathrm{Av}(E; \mathcal{M})$ does not fork over E. It remains to show that $\mathrm{Av}(E; E)$ is stationary. For this, it suffices to show that if $\overline{a}$ realizes $\mathrm{Av}(E, E)$ then for any finite equivalence relation R over E, for some $\overline{e} \in E$, $R(\overline{x}; \overline{e}) \in t(\overline{a}; E)$. But this is immediate from the definition of average type.

The next exercise illustrates another aspect of the analogy between indiscernible sets and models.

2.8 Exercise. Show directly from the definition that if $E \subset B$ is an infinite set of indiscernibles then $\mathrm{Av}(E, B)$ is finitely satisfied in E and thus $\mathrm{Av}(E, B)$ does not fork over E.

2.9 Exercise. Let M be $|A|^+$-saturated and A a subset of M. Show that if $p \in S(M)$ does not fork over A and the sequence $\langle \overline{e}_n : n < \omega \rangle$ is chosen so that $E_0 = A$ and $\overline{e}_{n+1} \in M$ realizes $p|E_n$ then $p = \mathrm{Av}(E; M)$.

With this in hand, we see that $\mathrm{Av}(E; E)$ is the desired syntactic description of E. The next definition makes more precise the requirement we would like this syntactic definition to meet. We proceed by formalizing the solution to the first problem of redundancy discussed at the beginning of this section.

2.10 Definition. Let p be a type over A and q a type over B. The types p and q are *parallel*, written $p \parallel q$, if there is a global type $\hat{r}$ extending $p \cup q$ which does not fork over either $\mathrm{dom}\, p$ or $\mathrm{dom}\, q$.

If p is stationary, p has a unique nonforking extension to a global type. Thus, parallelism is an equivalence relation on stationary types. In fact, parallelism is a congruence relation with respect to all the major relations among stationary types and this will be one of the first properties we prove as we introduce each new notion.

We could have replaced 'global type' in the definition of parallel by 'type over some model containing $\mathrm{dom}\, p \cup \mathrm{dom}\, q$'. This is easily verified using transitivity of independence in one direction and monotonicity in the other.

2.11 Theorem. *Suppose E and F are infinite sets of indiscernibles. Then $E \sim F$ if and only if* $\mathrm{Av}(E; E) \parallel \mathrm{Av}(F; F)$.

Proof. Let p (resp. q) be the unique global type which is a nonforking extension of $\mathrm{Av}(E; E)$ (resp. $\mathrm{Av}(F; F)$). If $p = q$, it is easy to restrict p and q to some sufficiently saturated model M containing $E \cup F$ and then define a strongly independent sequence realizing $p|M$ which will witness the equivalence of E and F.

Conversely, if E_0 witnesses the equivalence of E and F and r is the unique nonforking extension of $\mathrm{Av}(E_0, E_0)$ to a global type, it is easy to see that $p = r = q$.

In general we can restrict the size of the set E over which we average by invoking Theorem 1.23 and the following easy remark.

2.12 Lemma. *If E is a set of indiscernibles over A and an infinite subset of E is strongly independent over $A \cup E_0$ (for some subset E_0 of E) then E is strongly based on $E_0 \cup A$.*

2.13 Exercise. Let T be the theory of an equivalence relation E with two classes. Let p be the unique type over the empty set. Suppose $\neg E(a, b)$, p_a, p_b are non-forking extensions of p in $S(a)$, $S(b)$. Show $p_a \parallel p$, $p_b \parallel p$ but p_a is not parallel to p_b.

2.14 Historical Notes. This discussion of the average of a type is primarily taken from Chapter III of [Shelah 1978]. There are two complementary techniques for discussing the eventual behavior of a stationary type. Shelah builds a set of indiscernibles based on this type and works directly with this set of indiscernibles. Lascar emphasizes the role of the base type. We have incorporated both viewpoints; this section contains the fundamental properties of the indiscernible sequences. The definition of dimension appears in Chapter III of [Shelah 1978].

3. Forking and Dividing

In this section we study the combinatorial core of forking. We have defined the notion of forking in terms of definability of types in a stable theory. The concept we call implicit division here is Shelah's original notion of forking, which is described by specific properties of formulas. This formulation makes sense in an arbitrary theory and there have been some uses of it in unstable theories [Shelah 1980a]. We will not rely on results from this chapter except for examples. It is useful for reading Shelah's book and it does make the notion of forking much more concrete. We could have introduced implicit division at any time, but the proof of the equivalence of this notion with the one we introduced in Chapter III requires some of the properties of indiscernibles discussed in Section V.1.

We begin by defining some combinatorial properties of formulas and types. Note that types may be incomplete unless they are explicitly asserted to be complete.

3.1 Definition. For $n > 1$ and ϕ an L-formula, the collection of formulas $\{\phi(\overline{x}; \overline{e}) : \overline{e} \in E\}$ is *n-inconsistent* if for any subset E_0 of E with n or more elements, $\{\phi(\overline{x}; \overline{e}) : \overline{e} \in E_0\}$ is inconsistent but for each $\overline{e}$, $\models (\exists \overline{x})\phi(\overline{x}; \overline{e})$.

Thus, 2-inconsistency means the sets defined by the formulas $\phi(\overline{x}; \overline{e})$ are non-empty and disjoint for distinct choices of $\overline{e}$. For a slightly more complicated example, consider the set E of natural numbers which are divisible by at most 7 distinct primes. Define a single binary relation $R(x, y)$ which holds of n and p just if p is a prime and $p|n$. Then $\{R(x, e) : e \in E\}$ is 8-inconsistent.

3.2 Definition. The formula $\phi(\overline{x}; \overline{y})$ *divides* over A with respect to the type $q(\overline{y}) \in S(A)$ if i) $(\exists \overline{x})\phi(\overline{x}; \overline{y})$ is in $q(\overline{y})$ and ii) there exists an infinite set E of sequences which realize q such that $\{\phi(\overline{x}, \overline{e}) : \overline{e} \in E\}$ is n-inconsistent for some n.

We frequently write $\phi(\overline{x}; \overline{c})$ divides over A to mean, '$\phi(\overline{x}; \overline{y})$ divides over A with respect to $t(\overline{c}; A)$.' We may also write '$\phi(\overline{x}; \overline{y})$ and q divide over A' to indicate that ϕ divides with respect to q. However, the assertion, 'p divides over A' means there is a formula $\phi(\overline{x}; \overline{c}) \in p$ such that $\phi(\overline{x}; \overline{y})$ divides over A with respect to $t(\overline{c}; A)$. Note that if $t(\overline{c}; A)$ is algebraic, in particular if $\overline{c} \in A$, then $\phi(\overline{x}; \overline{c})$ does not divide over A.

We would like to use the notion of dividing as our combinatorial equivalent of forking and eventually for stable theories we will be able to do so (see Lemma 3.9). For the moment, however, we introduce a notion of implicit dividing and show it is equivalent to forking. We must do this as it is not obvious that if p is a type over B and p does not divide over A, then p can be extended to a complete type over B which does not divide over A.

3.3 Definition. The type $p(\overline{x})$ *implicitly divides* over A, if there exist formulas $\phi_0(\overline{x}; \overline{a}_0), \ldots, \phi_n(\overline{x}; \overline{a}_n)$ such that p implies the disjunction of the ϕ_i and each ϕ_i divides over A.

There is no assumption that the $\overline{a}_i$ realize the same type over A or even have the same length. We want first to establish some simple properties of implicit division and then show that it satisfies the extension requirement.

3.4 Lemma. i) *If $\overline{y}$ is a subsequence of $\overline{x}$, $\phi(\overline{x}; \overline{c})$ divides over A, and $\models (\exists \overline{y})\psi(\overline{y}; \overline{c})$, $\models \psi(\overline{y}; \overline{c}) \rightarrow \phi(\overline{x}; \overline{c})$, then $\psi(\overline{y}; \overline{c})$ divides over A.*

ii) *If $\phi(\overline{x}; \overline{a})$ divides over A then $\phi(\overline{x}; \overline{a})$ implicitly divides over A.*

iii) *The type p implicitly divides over A iff some finite subtype of p implicitly divides over A.*

iv) *(monotonicity) If A is contained in B and p, q are types such that q implies p (in particular if q extends p) then if p implicitly divides over B it follows that q implicitly divides over A.*

v) *If for each $i < n$, $p \cup \{\phi_i\}$ implicitly divides over A then $p \cup \bigvee_{i<n} \phi_i$ implicitly divides over A.*

Proof. All of these results follow immediately from the definitions.

We now consider a slightly more complicated condition which is equivalent to dividing. The proof is just a variant on the Ehrenfeucht-Mostowski method of constructing indiscernibles; such variants will abound in this section.

3.5 Lemma. *The formula $\phi(\overline{x};\overline{c})$ divides over A iff there is an infinite set $E = \{\overline{e}_i : i \in I\}$ of indiscernibles over A such that $\{\phi(\overline{x};\overline{e}) : \overline{e} \in E\}$ is n-inconsistent for some n and $\overline{c} = \overline{e}_0$.*

Proof. Clearly the latter condition implies that $\phi(\overline{x};\overline{c})$ divides over A. For the converse, suppose that F is a set of sequences such that $\{\phi(\overline{x};\overline{f}) : \overline{f} \in F\}$ is n-inconsistent and for each $\overline{f} \in F$, $t(\overline{f};A) = t(\overline{c};A)$. By Ramsey's theorem and compactness, there is a set E of indiscernibles over A such that each member of E realizes $t(\overline{c};A)$ and $\{\phi(\overline{x};\overline{e}_i) : \overline{e}_i \in E\}$ is n-inconsistent. Now choose an automorphism of the monster model which maps $\overline{e}_0$ to $\overline{c}$. The image of E under this automorphism satisfies the lemma.

Now by Lemma 1.15 we can rephrase Lemma 3.5 as follows.

3.6 Theorem. *If $\phi(\overline{x};\overline{y})$ and $t(\overline{c};A)$ divide over A, then there is an infinite set E with $\overline{c} \in E$ such that*

i) *any two sequences of the same length from E satisfy the same strong type over A and*
ii) *$\{\phi(\overline{x};\overline{e}) : \overline{e} \in E\}$ is n-inconsistent.*

Our next step is to show that for a stable theory, forking is the same as implicitly dividing.

3.7 Lemma. *If p implicitly divides over A then p forks over A.*

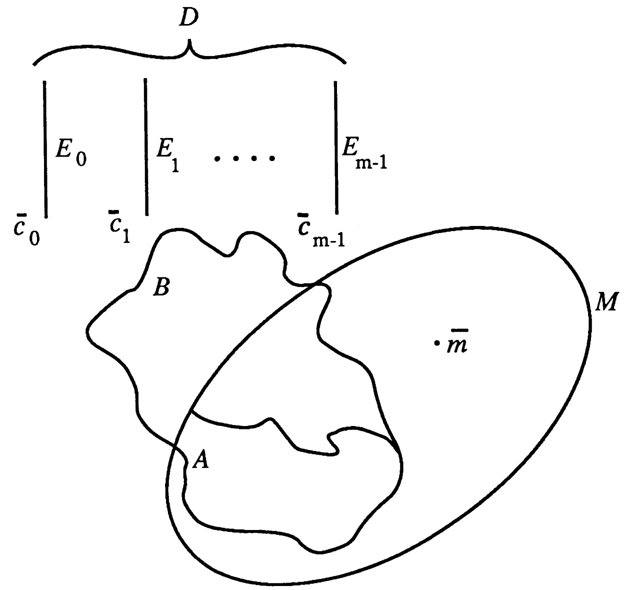

Fig. 3. Lemma V.3.7.

Proof. (Fig. 3). Let $B = \operatorname{dom} p$. If p does not fork over A, we can extend p to p' in $S(B)$ such that p' does not fork over A. If p implicitly divides over A, then by Lemma 3.4 iv) p' implicitly divides over A. Thus, if there is any counterexample p to the theorem, we may without loss of generality assume that $p \in S(B)$.

Now p implies a disjunction of, say m, formulas $\phi_i(\overline{x}; \overline{c}_i)$ which divide over A. That is, for each $i < m$ there is a set E_i of indiscernibles over A with $\overline{c}_i$ in E_i such that $\{\phi_i(\overline{x}, \overline{c}) : \overline{c} \in E_i\}$ is n-inconsistent for some n. Let D denote the union of the E_i. Choose M containing A such that $(M \downarrow D; A)$. Now p is definable over M (by Theorem III.3.26) since p does not fork over A. So we can extend p to $p_1 \in S(M \cup B \cup D)$ which is definable over M. Then for each formula $\phi(\overline{x}; \overline{y})$, there is a formula $d\phi(\overline{y}; \overline{m})$ in $F(M)$ such that for $\overline{e} \in M \cup B \cup D$, $\phi(\overline{x}; \overline{e})$ is in p_1 iff $d\phi(\overline{e}; \overline{m})$. By Lemma 1.18, since $t(\overline{m}; A \cup D)$ does not fork over A, E_i is a set of indiscernibles over $A \cup \{\overline{m}\}$. Now p implies $\bigvee_{i<m} \phi_i(\overline{x}; \overline{c}_i)$ and p is contained in the complete type p_1 so, for some $\phi_i(\overline{x}; \overline{c}_i)$, call it $\phi(\overline{x}, \overline{c})$, $\phi(\overline{x}, \overline{c}) \in p_1$, so $d\phi(\overline{c}; \overline{m})$ holds. But then, by indiscernibility, $d\phi(\overline{e}; \overline{m})$ holds for all $\overline{e}$ in E_i. Thus $\phi(\overline{x}, \overline{e})$ is in p_1 for all $\overline{e}$ in E_i, contradicting n-inconsistency and proving the lemma.

The next lemma justifies the extension from divides to implicitly divides. Note that it does not require stability.

3.8 Lemma. *If p is a type over B which does not implicitly divide over A, then p can be extended to a complete type q in $S(B)$ which does not implicitly divide over A.*

Proof. Let $W = \{\phi(\overline{x}) \in F(B) : \phi$ implicitly divides over $A\}$. Let q' be $p \cup \{\neg\phi(\overline{x}) : \phi \in W\}$. We first show q' is consistent. If not, there is a finite subset p' of p and a finite subset W_0 of W such that $p' \cup \{\neg\phi : \phi \in W_0\}$ is inconsistent. But then p implies the disjunction of the formulas $\phi(\overline{x})$, in W_0. Since each $\phi(\overline{x})$ implicitly divides over A, by Lemma 3.4 iv) and v) p implicitly divides over A.

Let q be any extension of q' in $S(B)$; we now show that q does not implicitly divide over A. If it does, then q implies the disjunction of a finite number of formulas, each of which implicitly divides over A. Then some finite subset, q_0, of q implies this disjunction so if $\theta(\overline{x})$ denotes the conjunction of q_0, $\theta(\overline{x})$ implicitly divides over A. So $\neg\theta$ is in q' which is contained in q, contradiction.

We have shown that 'implicit dividing' implies 'forking'. We now show a result somewhat stronger than the converse. Namely we show 'forking' implies 'dividing' (for a stable theory). The original proof of the following result was considerably shortened by a suggestion of Jürgen Saffe.

3.9 Lemma. *Let T be a stable theory $A \subseteq B$, and p a type over B which is closed under finite conjunction. If p forks over A then p divides over A.*

Proof. If p forks over A then by Corollary III.3.13, there is a formula $\phi(\overline{x}; \overline{b}) \in p$ such that $\{\phi(\overline{x}; \overline{b}')\}$ forks over A for any $\overline{b}'$ realizing $q = t(\overline{b}; A)$. It suffices to show that $\phi(\overline{x}; \overline{b})$ and q divides over A. Note that $t(\overline{b}; A)$ cannot be algebraic, since if it were $\phi(\overline{x}; \overline{b})$ would be almost over A which would contradict Corollary IV.1.5. Let E, with $|E| \geq \kappa(T)$, be an independent set of sequences realizing q. Choose E to be strongly based on A. If $\phi(\overline{x};\overline{b})$ and q do not divide over A then $\{\phi(\overline{x};\overline{e}) : \overline{e} \in E\}$ is n-consistent for every n. But then for some $\overline{c} \in \mathcal{M}$, $\models \phi(\overline{c};\overline{e})$ for all $\overline{e} \in E$. But, $\overline{c} \not\downarrow_A \overline{e}$ for each $\overline{e} \in E$, which contradicts Theorem II.2.18.

The following example shows the necessity both here and in Corollary III.3.13 of assuming that the type p is closed under conjunction.

3.10 Example. Let the language L contain three unary relation symbols, U_1, U_2, U_3 and a binary relation symbol E. The theory T asserts that the U_i partition the universe into infinite sets and that E is an equivalence relation splitting U_2 into infinitely many infinite classes but each element of U_1 and U_3 is related only to itself. Let b be an element of U_2. Then neither of the formulas, $\phi_1 \equiv E(x,b) \vee U_1(x)$, nor $\phi_2 \equiv E(x,b) \vee U_3(x)$ forks over the empty set but the formula $\phi_1(x,b) \wedge \phi_2(x,b)$ does.

Even without assuming that p is closed under conjunction we easily obtain the following.

3.11 Lemma. *If p is a type over B and p forks over A then p implicitly divides over A.*

Proof. If p does not implicitly divide over A we can, by Lemma 3.8, extend p to a complete type over B which does not implicitly divide over A. By the first monotonicity axiom this complete type forks over A so without loss of generality we may assume we have a counterexample p, to the lemma which is in $S(B)$. Now, by Lemma 3.9 we have the result.

The following example due to Charles Steinhorn shows the necessity of assuming T is stable for the results in this section.

3.12 Example. We construct a theory T and a formula ϕ such that ϕ implicitly divides over the empty set but ϕ does not divide, nor even fork, over the empty set. Let T have a language with two unary relations U_1 and U_2 and one binary relation R. The prototypical model of T interprets U_1 as an infinite set X, U_2 as the power set of X, and R as the membership relation. Suppose a and b are elements of U_2 such that a names the complement of the set named by b. Then the formula $R(x,a) \vee R(x,b)$ does not divide over the empty set since it is equivalent to $U_1(x)$. But each of the disjuncts divides.

We can summarize this chapter as follows.

3.13 Corollary. *Let T be a stable theory, $A \subseteq B$ and $p \in S(B)$. The following are equivalent.*

i) *p forks over A.*
ii) *p divides over A.*
iii) *p implicitly divides over A.*

3.14 Historical Notes. The term 'implicitly divides' is new here. The equivalence of dividing and implicitly dividing for stable theories was first shown by Shelah. Further work on dividing can be found in [Shelah 1980a].