CHAPTER 2
CLASSIFICATIONS

In this Chapter, we treat several significant BRT fragments. For most of these BRT fragments, we show that every statement is either provable or refutable in RCA_0.

For the remainder of these BRT fragments, we show that every statement is either provable in RCA_0, refutable in RCA_0, or provably equivalent to the Thin Set Theorem of section 1.4 over RCA_0.

Thus in this Chapter, we do not run into any independence results from ZFC. In the classification of Chapter 3, we do run into a statement independent of ZFC, called the Principal Exotic Case, which is the focus of the remainder of the book.

In this Chapter, we focus on five BRT settings (see Definition 1.1.11). These fall naturally, in terms of their observed BRT behavior, into three groups (see Definitions 1.1.2, and 2.1):

(SD,INF), (ELG ∩ SD,INF).
(ELG,INF), (EVSD,INF).
(MF,INF).

The inclusion diagram for these five sets of multivariate functions is

ELG ∩ SD
SD ELG
EVSD
MF

Here each item in any row is properly contained in any item in any lower row. Multiple items on any row are incomparable under inclusion.

(SD,INF), (ELG,INF), and (MF,INF) are the most natural of these five BRT settings. The remaining two BRT settings are closely associated, and serve to round out the theory.

MF (multivariate functions), SD (strictly dominating), and INF (infinite) were defined in section 1.1 in connection with the Complementation Theorem and the Thin Set Theorem.

DEFINITION 2.1. Let f ∈ MF. We say that f is of expansive linear growth if and only if there exist rational constants c,d > 1 such that for all but finitely many x ∈ dom(f),

$$c|x| \leq f(x) \leq d|x|$$

where |x| is the maximum coordinate of the tuple x. Let ELG be the set of all f ∈ MF of expansive linear growth.

DEFINITION 2.2. Let f ∈ MF. We say that f is eventually strictly dominating if and only if for all but finitely many x ∈ dom(f), f(x) > |x|. We write EVSD for the set of all f ∈ MF that are eventually strictly dominating.

In this Chapter, the two asymptotic BRT settings (ELG,INF), (EVSD,INF), have the same behavior, whereas the two non asymptotic BRT settings (SD,INF), (ELG ∩ SD,INF), also have the same behavior. In this Chapter, the behavior of (ELG,INF), (EVSD,INF) differs from the behavior of (SD,INF), (ELG ∩ SD,INF). In this Chapter, (MF,INF) behaves differently from the other four settings.

2.1. Methodology.

In this section, we use notation and terminology that was introduced in section 1.1.

Recall the definitions of

BRT fragment. Definition 1.1.18.
BRT environment. Definition 1.1.19.
BRT signature. Definition 1.1.21.
flat BRT fragment. Definition 1.1.34.

In Definition 1.1.39, the flat BRT fragments were divided into these four mutually disjoint categories:

1) EBRT in σ on (V,K), where σ does not end with $\subseteq$.
2) EBRT in σ on (V,K), where σ ends with $\subseteq$.
3) IBRT in σ on (V,K), where σ does not end with $\subseteq$.
4) IBRT in σ on (V,K), where σ ends with $\subseteq$.

Let α be a flat BRT fragment, and let S be an α format; i.e., a set of α elementary inclusions. According to Definition 1.39, we say that S is α correct if and only if

1’) $(\forall g_1,\ldots,g_n \in V)(\exists B_1,\ldots,B_m \in K)(S)$.
2’) $(\forall g_1,\ldots,g_n \in V)(\exists B_1 \subseteq \ldots \subseteq B_m \in K)(S)$.
3’) $(\exists g_1,\ldots,g_n \in V)(\forall B_1,\ldots,B_m \in K)(S)$.
4’) $(\exists g_1,\ldots,g_n \in V)(\forall B_1 \subseteq \ldots \subseteq B_m \in K)(S)$.

where we use 1’),2’),3’),4’) according to whether α is in category 1),2),3),4).

For example, the Thin Set Theorem is the negation of a statement of the form 3').

In the case of EBRT and IBRT in A,fA on any given setting, there are 16 formats, and hence 16 statements of forms 1',3', respectively, that have to be considered. This is such a small number that we can profitably list all of these statements, and determine their truth values. We do this in section 2.2.

In the case of EBRT and IBRT on A,fA,fU on any given setting, there are 256 formats, and hence at most 256 statements that have to be considered. Actually, a closer look shows that there are only 6 elementary inclusions, generating only $2^6 = 64$ formats. In section 2.3, we list these formats in order of increasing cardinality. This avoids considerable duplication of work. This method of compilation is seen to be perfectly manageable in section 2.3.

In the case of EBRT and IBRT on A,B,fA,fB, there are $2^{16} = 65,536$ formats, and hence 65,536 statements that have to be considered. We do not attempt to work with A,B,fA,fB here.

In sections 2.4 - 2.7, we instead work with A,B,fA,fB,$\subseteq$. There are 9 elementary inclusions, and so $2^9 = 512$ formats need be considered. This is considerably less than 65,536.

Here a treelike methodology is preferable to the enumeration procedure used in section 2.3. We expect the treelike methodology to be the method of choice when analyzing richer BRT fragments.

We treat EBRT in A,B,fA,fB,$\subseteq$ on (SD,INF), (ELG,INF) in sections 2.4, 2.5. We treat EBRT in $A_1,\ldots,A_k,fA_1,\ldots,fA_k,\subseteq$ on (MF,INF) in section 2.6. We treat IBRT in $A_1,\ldots,A_k,fA_1,\ldots,fA_k,\subseteq$ on (SD,INF), (ELG,INF), (MF,INF) in section 2.7.

The most substantial uses of the treelike methodology are in sections 2.4 and 2.5. We believe that EBRT in A,B,fA,fB on (MF,INF), (SD,INF), (ELG,INF) can be treated using this treelike methodology, but with considerably more effort.

In this section, we rigorously present this treelike methodology and establish some important facts about it.

Fix a flat BRT fragment α. Let S be an α format. Let $1 \leq i \leq 4$ be such that α is in category i) above. We also fix a true formal system T that includes RCA_0. We assume that α is given a description in T.

According to Definition 1.1.42, we say that S is α,T correct if and only if T proves that S is α correct. We say that S is α,T incorrect if and only if T refutes that S is α correct.

According to Definition 1.43, we say that α is T secure if and only if every α format is α,T correct or α,T incorrect.

The goal of the treelike methodology is

a) to show that α is T secure.
b) to list all maximal α,T correct formats; i.e., α,T correct α formats that are not properly included in any α,T correct α format.

Note the following obvious but crucial property of α,T correct/incorrect α formats:

Every subset of an α,T correct α format is α,T correct. Every α format that contains an α,T incorrect format is α,T incorrect.

Goal b) is preferable to listing all α,T correct α formats, as the latter may be uncomfortably numerous, or even impractically enormous, whereas the former may be very manageable in size.

The challenge is to show that our treelike methodology does in fact rigorously justify the claim that we have actually established a) and done b). In other words, we need to justify that

i. The α formats listed under b) are indeed α,T correct, and are incomparable under inclusion.
ii. Any α format not included in any of those listed under b) is α,T incorrect.

Some readers may be content with examining the classifications made in sections 2.4, 2.5, and absorbing the methodology from the displays. When the significance of some features are not apparent, the reader can look at the formal treatment of the methodology presented below.

Let α be a flat BRT fragment with BRT setting (V,K). Recall the definition of α formulas (Definition 1.1.25).

DEFINITION 2.1.1. We say that an α formula is α,T valid if and only if, it is provable in T that it holds for all values of the function variables from V and all values of the set variables from K. In case the signature of α ends with $\subseteq$, the values of the set variables, in increasing order of subscripts, are assumed to form a tower under inclusion.

DEFINITION 2.1.2. An α worklist is a two part finite sequence

$$(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$$

where $r,s \geq 0$, and $\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s$ are α inclusions.

DEFINITION 2.1.3. The formats of an α worklist are the α formats that include $\{\varphi_1,\ldots,\varphi_r\}$ and are included in $\{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s)$.

DEFINITION 2.1.4. We say that a worklist $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is α,T secure if and only if for all $\{\varphi_1,\ldots,\varphi_r\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$, S is α,T correct or α,T incorrect.

Informally, the goal of an α worklist is to constructively verify that it is α,T secure, in the sense of determining the α,T correctness or α,T incorrectness of all α formats.

Sometimes we want to replace one worklist with a simpler worklist, without altering its goal. Here are some reduction operations that are very useful.

Let W = $\{\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$.

i. We can replace $\varphi_1,\ldots,\varphi_r$ with any $\varphi_1',\ldots,\varphi_p'$ such that $(\varphi_1 \wedge \ldots \wedge \varphi_r) \leftrightarrow (\varphi_1' \wedge \ldots \varphi_p')$ is α,T valid.

ii. We can replace any ψ_i by ψ_i', where $(\varphi_1 \wedge \ldots \wedge \varphi_r) \rightarrow (\psi_i \leftrightarrow \psi_i')$ is α,T valid.

iii. We can remove any ψ_i such that $(\varphi_1 \wedge \ldots \wedge \varphi_r) \rightarrow \psi_i$ is α,T valid.

iv. We can remove any ψ_i such that $(\varphi_1,\ldots,\varphi_r;\psi_i)$ is α,T-incorrect.

v. We can remove duplicates among $\psi_1,\ldots,\psi_s$.

vi. We can permute the $\psi_1,\ldots,\psi_s$.

DEFINITION 2.1.5. α,T reduction consists of performing any finite number of the above operations in succession.

This notion of α,T reduction corresponds to what happens in the classifications in sections 2.4, 2.5. For instance, consider LIST 1.2.1 in section 2.4.

Here the BRT fragment α is EBRT in A,B,fA,fB,$\subseteq$ on (SD,INF), and T is RCA_0. This displays the worklist (A $\cap$ fA = $\emptyset$, A $\cap$ fB = $\emptyset$, fA $\subseteq$ B; B $\cup$ fB = N, B $\subseteq$ A $\cup$ fB, fB $\subseteq$ B $\cup$ fA, B $\cap$ fB $\subseteq$ A $\cup$ fA). This gets reduced to the worklist displayed by LIST 1.2.1.*, which is the worklist (A $\cap$ fA = $\emptyset$, A $\cap$ fB = $\emptyset$, fA $\subseteq$ B; B $\cup$ fB = N, B $\subseteq$ A $\cup$ fB, B $\cap$ fB $\subseteq$ fA).

Here we have merely eliminated fB $\subseteq$ B $\cup$ fA from the second half of LIST 1.2.1, since Lemma 2.4.4 tells us that (A $\cap$ fB = $\emptyset$; fB $\subseteq$ B $\cup$ fA) is α,T incorrect.

LEMMA 2.1.1. Let α be a flat BRT fragment, and T be a true theory with a presentation of α. Suppose W =

$(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ α,T reduces to W' = $(\varphi_1',...,\varphi_p';\psi_1',...,\psi_q')$. Then W is α,T secure if and only if W' is α,T secure.

Proof: It suffices to show that α,T security is preserved under each of the operations i-vi. Let one of the operations send worklist W to worklist W'. In cases i,ii,iii,v,vi, evidently every α format for W is α,T equivalent to some α format for W', and vice versa.

It remains to consider operation iv. We have W = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$, W' = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_{i-1},\psi_{i+1},...,\psi_s)$, where W' is α,T secure. Let $\{\varphi_1,...,\varphi_r\} \subseteq S \subseteq \{\varphi_1,...,\varphi_r,\psi_1,...,\psi_s\}$. If $\psi_i \in S$ then S is α,T incorrect. If $\psi_i \notin S$ then S is a format for W_i', and so S is α,T correct or α,T incorrect. QED

It is simpler to use sequences instead of sets. Accordingly, let W = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ be an α worklist.

DEFINITION 2.1.6. A subsequence for W is a subsequence of the sequence $(\varphi_1,...,\varphi_r,\psi_1,...,\psi_s)$ that begins with $\varphi_1,...,\varphi_r$, and which includes the underlying subsequence of positions 1,...,r,...,r+s. This is very useful for handling all sorts of repetitions in worklists.

DEFINITION 2.1.7. A finite sequence of α elementary inclusions is said to be α,T correct (α,T incorrect) if and only if its set of terms is α,T-correct (α,T incorrect).

LEMMA 2.1.2. Let α be a flat BRT fragment, and T be a true theory with a presentation of α. Let an α,T reduction of W = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ to W' = $(\varphi_1',...,\varphi_p';\psi_1,...,\psi_q')$ be given. Let the list of maximal α,T correct subsequences for W' be given (together with proofs in T). We can efficiently generate the list of maximal α,T correct subsequences for W (together with proofs in T). Furthermore, these two lists have the same number of sequences.

Proof: We can assume that we have W = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ that is α,T reduced to W' = $(\varphi_1',...,\varphi_p';\psi_1',...,\psi_q')$ by any one of the reductions i-v.

case i. Here we have W = $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ and W' = $(\varphi_1',...,\varphi_p';\psi_1,...,\psi_s)$. Let f be the obvious one-one correspondence between subsequences for W and subsequences for W'. Then for every α sequence τ for W, τ and f(τ) are

α,T equivalent. It is now evident that τ is α,T correct if and only if f(τ) is α,T correct. It is then evident that τ is maximally α,T correct for W if and only if f(τ) is maximally α,T correct for W'.

case ii. Here we have W = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ and W' = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_{i-1},\psi_i',\psi_{i+1},\ldots,\psi_s)$. Let f be the obvious one-one correspondence between subsequences for W and subsequences for W', based on corresponding positions. Then for every α sequence τ for W, τ and f(τ) are α,T equivalent. As in case i, τ is maximally α,T correct for W if and only if f(τ) is maximally α,T correct for W'.

case iii. Here we have W = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$, W' = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_{i-1},\psi_{i+1},\ldots,\psi_s)$. Let f be the obvious map from subsequences for W to subsequences for W' defined by ignoring ψ_i; i.e., as position r+i. Note that f is not one-one. However, the restriction g of f to the τ with ψ_i (as position r+i) is one-one, and for all $\tau \in$ dom(g), τ and g(τ) are α,T equivalent. Since $(\varphi_1 \wedge \ldots \wedge \varphi_r) \rightarrow \psi_i$ is α,T valid, all maximal α,T correct subsequences for W have ψ_i (as position r+i). It is now evident that g is a one-one correspondence between the maximal α,T correct subsequences for W and the maximal α,T correct subsequences for W'.

case iv. Here we have W = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$, W' = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_{i-1},\psi_{i+1},\ldots,\psi_s)$. Note that the α,T correct subsequences for W are identical to the α,T correct subsequences for W', since ψ_i cannot be present.

cases v-vi. Left to the reader.

QED

LEMMA 2.1.3. Let α be a flat BRT fragment, and T be a true theory with a presentation of α. Let an α,T reduction of W = $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ to W' = $(\varphi_1',\ldots,\varphi_p';\psi_1,\ldots,\psi_q')$ be given. Let the list of maximal α,T correct formats for W' be given (together with proofs in T). We can efficiently generate the list of maximal α,T correct formats for W (together with proofs in T). Furthermore, these two lists have the same number of formats. W is α,T secure if and only if W' is α,T secure.

Proof: This is the same as Lemma 2.1.2, except that we are using subsets (formats) instead of subsequences. It suffices to observe that the maximal α,T correct sequences for W are exactly the subsequences for W whose set of terms is an α,T correct format for W. The last claim is by Lemma 2.1.2. QED

T CLASSIFICATIONS FOR BRT FRAGMENTS

DEFINITION 2.1.8. The starred α worklists are the α worklists with a * appended at the end.

DEFINITION 2.1.9. We say that TREE is a T classification for α if and only if α is a flat BRT fragment, T is a true theory extending RCA_0 which adequately defines the BRT setting of α, and TREE is a finite labeled tree with the properties given below.

1. The root of TREE is labeled by an α worklist $(;\delta_1,...,\delta_t)$, where the δ's list all α elementary inclusions without repetition.

2. Suppose a vertex v is labeled $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$, where v is not terminal. Then v has exactly one son w. The label of w is some $(\varphi_1',...,\varphi_p';\psi_1',...,\psi_q')$*, where

$(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ is α,T reducible to $(\varphi_1',...,\varphi_p';\psi_1',...,\psi_q')$.
$\varphi_1',...,\varphi_p',\psi_1',...,\psi_q'$ are distinct.

In sections 2.4, 2.5, note that the worklists whose names don't end with * are immediately followed by those which do, and the succeeding worklists with * are obtained by α,T reduction.

3. Suppose a vertex v is labeled $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$*, where v is not terminal. Then there exists $1 \leq i \leq s$ such that v has exactly i sons $w_1,...,w_i$, with labels

$$(\varphi_1,...,\varphi_r,\psi_1;\psi_2,...,\psi_s)$$
$$(\varphi_1,...,\varphi_r,\psi_2;\psi_3,...,\psi_s)$$
$$...$$
$$(\varphi_1,...,\varphi_r,\psi_i;\psi_{i+1},...,\psi_s)$$

respectively, where w_i is terminal, and $w_1,...,w_{i-1}$ are not terminal.

4. Suppose the vertex v is terminal, with label $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$ or $(\varphi_1,...,\varphi_r;\psi_1,...,\psi_s)$*. Then $\{\varphi_1,...,\varphi_r,\psi_1,...,\psi_s\}$ is α,T correct.

This completes Definition 2.1.7.

We want to show that if we have a T classification for α, then α is T secure.

LEMMA 2.1.4. Let TREE be a T classification for α. Then α is T secure and the number of maximally α,T correct α formats is at most the number of terminal vertices of T.

Proof: We prove the following by induction on TREE. Let v be a vertex of TREE whose label is the worklist W (or W*). Then W is α,T secure, and the number of maximal α,T correct formats for W is the number of terminal vertices from v; i.e., the number of terminal vertices that descend from v, including v.

case 1. v is a terminal vertex of TREE. Let the label of v be $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ or $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)^*$. Then $(\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s)$ is α,T correct. Hence $\{\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is α,T secure. The number of maximal α,T correct formats for $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is 1.

case 2. Suppose v has label $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$, and is nonterminal. Then v has exactly one son, w, labeled $(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q')^*$. Suppose $(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q')$ is α,T secure. Suppose the number of maximal α,T correct formats for $(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q')$ is at most the number of terminal vertices from w. The label of w is some $(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q')^*$, where

$$(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q') \text{ is an } \alpha\text{,T reduction of } (\varphi_1,\ldots,\varphi_r;\varphi_1,\ldots,\varphi_s).$$

By the induction hypothesis, $(\varphi_1',\ldots,\varphi_p';\psi_1',\ldots,\psi_q')$ is α,T secure. Hence by Lemma 2.1.3, $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is α,T secure. Also by Lemma 2.1.3, the number of maximal α,T correct formats is preserved.

case 3. Suppose v has label $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)^*$, where v is not terminal. Let $1 \leq i \leq s$, where v has exactly i sons, $w_1,\ldots,w_i$, with labels

$$(\varphi_1,\ldots,\varphi_r,\psi_1;\psi_2,\ldots,\psi_s)$$
$$(\varphi_1,\ldots,\varphi_r,\psi_2;\psi_3,\ldots,\psi_s)$$
$$\ldots$$
$$(\varphi_1,\ldots,\varphi_r,\psi_i;\psi_{i+1},\ldots,\psi_s)$$

respectively, where w_i is terminal. Suppose each of these labels is α,T secure. Suppose for each $1 \leq j \leq i$, the number of maximal α,T correct formats for $(\varphi_1,\ldots,\varphi_r,\psi_j;\psi_{j+1},\ldots,\psi_s)$ is the number of terminal vertices from w_j.

Note that $\{\varphi_1,\ldots,\varphi_r,\psi_i;\psi i+1,\ldots,\psi_s\}$ is α,T correct, and so automatically α,T secure. Also note that $\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s$ are distinct.

Let $\{\varphi_1,\ldots,\varphi_r\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$. Suppose first that $S \cap \{\psi_1,\ldots,\psi_i\} \neq \emptyset$. Let $1 \leq j \leq i$ be least such that $\psi_j \in S$. Then $\{\varphi_1,\ldots,\varphi_r,\psi_j\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_{j+1},\ldots,\psi_s\}$. By the induction hypothesis, $(\varphi_1,\ldots,\varphi_r,\psi_j;\psi_{j+1},\ldots,\psi_s)$ is α,T secure. Hence S is α,T correct or α,T incorrect.

Now suppose that $S \cap \{\psi_1,\ldots,\psi_i\} = \emptyset$. Then $S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_{i+1},\ldots,\psi_s\}$. Hence S is α,T correct.

Now let S be maximal so that $\{\varphi_1,\ldots,\varphi_r\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$ and S is α,T correct. Suppose first that $S \cap \{\psi_1,\ldots,\psi_i\} \neq \emptyset$. Let $1 \leq j \leq i$ be least such that $\psi_j \in S$. Then $\{\varphi_1,\ldots,\varphi_r,\psi_j\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_{j+1},\ldots,\psi_s\}$. In fact, S is maximal such that $\{\varphi_1,\ldots,\varphi_r,\psi_j\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_{j+1},\ldots,\psi_s\}$.

Now suppose that $S \cap \{\psi_1,\ldots,\psi_i\} = \emptyset$. Then $S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_{i+1},\ldots,\psi_s\}$. Hence $S = \{\varphi_1,\ldots,\varphi_r,\varphi_{i+1},\ldots,\psi_s\}$.

Hence the number of maximal α,T correct formats for $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is at most the sum over $1 \leq j \leq i$ of the number of maximal α,T correct formats for the label of w_j. By the induction hypothesis, the number of maximal α,T correct formats for the label of w_j is at most the number of terminal vertices from w_j. Hence the number of maximal α,T correct formats for $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ is at most the number of terminal vertices from v.

This concludes the induction argument. Now apply the result to the label of the root. QED

THEOREM 2.1.5. Let α be a flat BRT fragment, and T be a true theory with a presentation of α. Then α is T secure if and only if there is a T classification for α. Let TREE be a T classification for α. The number of maximally α,T correct α formats is at most the number of terminal vertices of T.

Proof: Let α,T be as given. By Lemma 2.1.4, we need only show that if α is T secure, then there is a T classification for α.

Assume α is T secure. We build TREE as follows. The construction will be such that any vertex whose label is starred is not terminal.

Create the root of T, with label $(;\delta_1,\ldots,\delta_r)$, where $\delta_1,\ldots,\delta_r$ is a listing, without repetition, of the α elementary inclusions.

Suppose we have constructed the vertex v of TREE with label $W = (\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$. If $\{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$ is α,T correct, then v is terminal. Otherwise, we apply α,T reductions iii,iv,v to W, as much as possible, as well as removing duplicates among $\varphi_1,\ldots,\varphi_r$. Let the result be the worklist W'. We create the single son w of v, with label W*. Clearly W' is not α,T correct.

Suppose we have constructed the vertex v of TREE with label $W^* = (\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)^*$. If $\{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$ is α,T correct then v is terminal. Suppose $\{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$ is not α,T correct. Then v is not α,T correct. Clearly $\{\varphi_1,\ldots,\varphi_r,\psi_s\}$ is α,T correct, since otherwise we could apply reduction operation iv to $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$, contrary to W* being a label of a vertex in TREE.

Let $2 \leq i \leq s$ be smallest such that $\{\varphi_1,\ldots,\varphi_r,\psi_i,\ldots,\psi_s\}$ is α,T correct. Create i sons with labels

$$(\varphi_1,\ldots,\varphi_r,\psi_1;\psi_2,\ldots,\psi_s)$$
$$(\varphi_1,\ldots,\varphi_r,\psi_2;\psi_3,\ldots,\psi_s)$$
$$\ldots$$
$$(\varphi_1,\ldots,\varphi_r,\psi_i;\psi_{i+1},\ldots,\psi_s)$$

respectively, where w_i is terminal. Vertices $w_1,\ldots,w_{i-1}$ are not terminal.

This construction must terminate since

a. The clause applying to non starred vertices that are not terminal, creates a single son whose label has the same number of entries to the right of the semicolon.

b. The clause applying to starred vertices v that are not terminal, creates sons $w_1,\ldots,w_i$, where for all j, the number of entries to the right of the label of w_j is less than the number of entries to the right of the label of v.

QED

THEOREM 2.1.6. Let α be a flat BRT fragment, and T be a true theory with a presentation of α. Let TREE be a T classification for α. We can efficiency list all of the maximal α,T correct formats.

Proof: Let α,T,TREE be as given. For each worklist for vertices in TREE, $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$, we construct a list of the maximal α,T correct formats S with $\{\varphi_1,\ldots,\varphi_r\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$. We do this by recursion, starting at the terminal vertices, towards the root, ending at the root. At terminal vertices, there is exactly one maximal α,T correct S. At nonterminal non starred vertices, apply the procedure from Lemma 2.1.3.

Now let $(\varphi_1,\ldots,\varphi_r;\psi_1,\ldots,\psi_s)$ be the worklist at a nonterminal starred vertex. Let $1 \leq i \leq s$ be such that the vertex has the i sons with labels

$$(\varphi_1,\ldots,\varphi_r,\psi_1;\psi_2,\ldots,\psi_s)$$
$$(\varphi_1,\ldots,\varphi_r,\psi_2;\psi_3,\ldots,\psi_s)$$
$$\ldots$$
$$(\varphi_1,\ldots,\varphi_r,\psi_i;\psi_{i+1},\ldots,\psi_s)$$

respectively, where w_i is terminal. Vertices $w_1,\ldots,w_{i-1}$ are not terminal.

We already have the i lists of maximal α formats associated with each of the above i worklists. Clearly every maximal α,T correct format S with $\{\varphi_1,\ldots,\varphi_r\} \subseteq S \subseteq \{\varphi_1,\ldots,\varphi_r,\psi_1,\ldots,\psi_s\}$ must appear in at least one of these lists. So we can simply merge these lists of α formats, and take their maximal elements. QED

The tree methodology we have presented here is applicable to situations that do not involve BRT.

An important application of this tree methodology occurs in section 2.7 (see Witness Set List), where we start with a list of sets $V_1,V_2,\ldots,V_k$, and we want to determine which subsets of $\{V_1,\ldots,V_k\}$ have nonempty intersection. Thus the notion of “correctness” of subsets of $\{V_1,...,V_k\}$ here is “having a nonempty intersection”.

But what takes the place of the notion of reduction used in case 2? In the application in section 2.7, we only use the elimination of terms, from the second part of a worklist, that is disjoint from the intersection of the terms from the first part of that worklist.

This rather pure form of our tree methodology is used to prove Theorems 2.7.25 - 2.7.27.