

CHAPTER 1

INTRODUCTION TO BRT

- 1.1. General Formulation.
- 1.2. Some BRT Settings.
- 1.3. Complementation Theorems.
- 1.4. Thin Set Theorems.

1.1. General Formulation.

Before presenting the precise formulation of Boolean Relation Theory (BRT), we give two examples of assertions in BRT that are of special importance for the theory.

DEFINITION 1.1.1. N is the set of all nonnegative integers. $A \setminus B = \{x: x \in A \wedge x \notin B\}$. For $x \in N^k$, we let $\max(x)$ be the maximum coordinate of x .

THIN SET THEOREM. Let $k \geq 1$ and $f: N^k \rightarrow N$. There exists an infinite set $A \subseteq N$ such that $f[A^k] \neq N$.

COMPLEMENTATION THEOREM. Let $k \geq 1$ and $f: N^k \rightarrow N$. Suppose that for all $x \in N^k$, $f(x) > \max(x)$. There exists an infinite set $A \subseteq N$ such that $f[A^k] = N \setminus A$.

These two theorems are assertions in BRT. In fact, the complementation theorem has the following sharper form.

COMPLEMENTATION THEOREM (with uniqueness). Let $k \geq 1$ and $f: N^k \rightarrow N$. Suppose that for all $x \in N^k$, $f(x) > \max(x)$. There exists a unique set $A \subseteq N$ such that $f[A^k] = N \setminus A$. Furthermore, A is infinite.

We will explore the Thin Set Theorem and the Complementation Theorem in sections 1.3, 1.4. At this point we analyze their logical structure.

DEFINITION 1.1.2. A multivariate function on N is a function whose domain is some N^k and whose range is a subset of N . A strictly dominating function on N is a multivariate function on N such that for all $x \in N^k$, $f(x) > \max(x)$. We define MF as the set of all multivariate functions on N , SD as the set of all strictly dominating functions on N , and INF as the set of all infinite subsets of N .

DEFINITION 1.1.3. Let $f \in MF$, where $\text{dom}(f) = N^k$. For $A \subseteq N$, we define $fA = f[A^k]$.

The notation fA is very convenient. It avoids the unnecessary use of explicit mention of arity or dimension. It is used throughout this book.

Using this notation, we can restate our two theorems as follows.

THIN SET THEOREM. For all $f \in MF$ there exists $A \in INF$ such that $fA \neq N$.

COMPLEMENTATION THEOREM. For all $f \in SD$ there exists $A \in INF$ such that $fA = N \setminus A$.

Note that in the Thin Set Theorem, we use the family of multivariate functions MF , and the family of sets INF . In the Complementation Theorem, we use the family of multivariate functions SD , and the family of sets INF .

In BRT terminology this will be expressed by saying that the Thin Set Theorem is an instance of IBRT (inequational BRT) on the BRT setting (MF, INF) , and the Complementation Theorem is an instance of EBRT (equational BRT) on the BRT setting (SD, INF) .

Note that we can regard the condition $fA \neq N$ as a Boolean inequation in fA, N . We also regard the condition $fA = N \setminus A$ as a Boolean equation in fA, N .

Here N plays the role of the universal set in Boolean algebra. From this perspective, $fA \neq N$ is a Boolean inequation in fA , and $fA = N \setminus A$ is a Boolean equation in A, fA .

The fact that N should play the role of the universal set can be read off from the BRT settings (MF, INF) and (SD, INF) . See "Full BRT Semantics" below.

EBRT stands for "equational Boolean relation theory". IBRT stands for "inequational Boolean relation theory".

Thus we say that

- i. The Thin Set Theorem is an instance of: IBRT in fA on (MF, INF) .

ii. The Complementation Theorem is an instance of: EBRT in A, fA on (SD, INF) .

We now fully explain what we mean by such phrases as "IBRT in fA on (MF, INF) " and "EBRT in A, fA on (SD, INF) ".

The principal BRT environments are

IBRT
EBRT

defined below. We will mention one other (much richer) BRT environment below (PBRT), but in this book we stay within the environments IBRT and EBRT.

We call the lists

fA
 A, fA

BRT signatures. In general, the BRT signatures will be substantially richer than the above two examples.

We have already called the pairs

(MF, INF)
 (SD, INF)

BRT settings. One other BRT setting plays a particularly important role in this book. This is the BRT setting (ELG, INF) . See Definition 2.1.

We are now prepared for the formal presentation of BRT.

FULL BRT SYNTAX

DEFINITION 1.1.4. The BRT set variables are the symbols A_1, A_2, \dots . The BRT function variables are the symbols f_1, f_2, \dots .

In practice, we will use appropriate upper case and lower case letters without subscripts for these BRT variables.

DEFINITION 1.1.5. The BRT terms are defined by

i) every BRT set variable is a term;

- ii) \emptyset, U are BRT terms (U represents the universal set);
- iii) if s, t are BRT terms then $(s \cup t), (s \cap t), (s \setminus t)$ are BRT terms;
- iv) if f is a BRT function variable and t is a BRT term then ft is a BRT term.

DEFINITION 1.1.6. The BRT equations are of the form $s = t$, where s, t are BRT terms. The BRT inequations are of the form $s \neq t$, where s, t are BRT terms. The BRT inclusions are of the form $s \subseteq t$, where s, t are BRT terms.

DEFINITION 1.1.7. The BRT formulas are defined by

- i) every BRT equation is a BRT formula;
- ii) if φ, ψ are BRT formulas then $(\neg\varphi), (\varphi \vee \psi), (\varphi \wedge \psi), (\varphi \rightarrow \psi), (\varphi \leftrightarrow \psi)$ are BRT formulas.

We routinely omit parentheses when no ambiguity arises. We also adhere to the usual precedence table

$$\begin{array}{c} \neg \\ \vee \wedge \\ \rightarrow \leftrightarrow \end{array}$$

FULL BRT SEMANTICS

DEFINITION 1.1.8. A multivariate function is a pair (f, k) , where

- i) f is a function in the standard sense of a univalent set of ordered pairs;
- ii) if $k \geq 2$, then every element of $\text{dom}(f)$ is a k -tuple.

DEFINITION 1.1.9. We say that the arity of (f, k) is k . The domain of (f, k) is taken to be $\text{dom}(f)$.

We rely on the fact that for all $1 < i < j$, no i -tuple is a j -tuple.

Let f be a function (in the standard sense). Note that if f is empty then for all $k \geq 1$, (f, k) is a multivariate function. Also, if f is nonempty then

- i) $(f, 1)$ is a multivariate function;
- ii) there is at most one $k \geq 2$ such that (f, k) is a multivariate function.

The explicit mention of k is intended to avoid the following type of ambiguity. A function $f:N^2 \rightarrow N$ could be viewed as either a 1-ary multivariate function with domain N^2 , or a 2-ary multivariate function with domain N^2 . In our notation, the former would be written $(f,1)$, and the latter would be written $(f,2)$. Note that $(f,3)$ is not a multivariate function.

In practice, the intended arity k of functions is clear from context, and we generally ignore the above definition of multivariate function. However, we need the above definition for full rigor.

DEFINITION 1.1.10. Let $f = (f,k)$ be a multivariate function and E be a set. We define $fE = f[E^k] = \{f(x_1, \dots, x_k) : x_1, \dots, x_k \in E\} = \{y : (\exists x_1, \dots, x_k \in E) (y = f(x_1, \dots, x_k))\}$.

DEFINITION 1.1.11. A BRT setting is a pair (V,K) , where V is a nonempty set of multivariate functions and K is a nonempty family of sets.

DEFINITION 1.1.12. The BRT assertions are the assertions of the form $(\forall g_1, \dots, g_n \in V) (\exists B_1, \dots, B_m \in K) (\varphi)$

where $n, m \geq 1$, B_1, \dots, B_m are distinct BRT set variables, g_1, \dots, g_n are distinct BRT function variables, and φ is a BRT formula involving at most the variables $B_1, \dots, B_m, g_1, \dots, g_n$.

Every BRT assertion gives rise to an actual mathematical statement provided we are also given a BRT setting (V,K) . Specifically,

DEFINITION 1.1.13. \cap is interpreted as intersection, \cup as union, \setminus as set theoretic difference, and \emptyset as the empty set. f_t is interpreted as the image of f on the interpretation of t , using Definition 1.10. U is interpreted as the least set U such that

- i) for all $A \in K$, $A \subseteq U$;
- ii) for all $f \in V$, $fU \subseteq U$.

Note that U may or may not lie in K .

An important special kind of BRT is obtained by requiring that the relevant sets form a tower under inclusion. Specifically,

DEFINITION 1.1.14. The BRT, \subseteq assertions are the assertions of the form

$$(\forall g_1, \dots, g_n \in V) (\exists B_1 \subseteq \dots \subseteq B_m \in K) (\varphi)$$

where $n, m \geq 1$, B_1, \dots, B_m are distinct BRT set variables, g_1, \dots, g_n are distinct BRT function variables, and φ is a BRT formula involving at most the variables $B_1, \dots, B_m, g_1, \dots, g_n$.

Here $B_1 \subseteq \dots \subseteq B_m \in K$ means

$$B_1 \subseteq \dots \subseteq B_m \wedge B_1, \dots, B_m \in K.$$

DEFINITION 1.1.15. We say that a BRT formula is BRT valid if and only if it is true on all BRT settings (V, K) under any assignment of elements of V to the function variables, and any assignment of elements of K to the set variables.

DEFINITION 1.1.16. We say that a BRT formula is BRT, \subseteq valid if and only if it is true on all BRT settings (V, K) under any assignment of elements of V to the function variables, and any assignment of elements of K to the set variables such that for all $i \leq j$, the assignment to A_i is a subset of the assignment to A_j .

DEFINITION 1.1.17. Let φ, ψ be BRT formulas. We say that φ, ψ are BRT (BRT, \subseteq) equivalent if and only if $\varphi \leftrightarrow \psi$ is BRT (BRT, \subseteq) valid. This definition is extended to sets of BRT formulas in the obvious way.

BRT FRAGMENTS

Obviously there are infinitely many BRT formulas. Results concerning all BRT formulas, even in very basic BRT settings, have been entirely inaccessible to us. The book will only be concerned with very modest fragments of BRT.

DEFINITION 1.1.18. The BRT fragments are written

$$[\text{Environment}] \text{ in } [\text{Signature}] \text{ on } [\text{Setting}].$$

It remains to say what the BRT Environments and Signatures are. The BRT Settings have already been defined.

DEFINITION 1.1.19. There are three BRT environments:

- i) EBRT (equational BRT);
- ii) IBRT (inequational BRT);
- iii) PBRT (propositional BRT).

DEFINITION 1.1.20. A core BRT term is a BRT term that is either a BRT set variable or begins with a BRT function variable. For example, $f_3(A_1 \cup A_4)$ is a core BRT term, and $A_1 \cup A_4$ is not a core BRT term.

DEFINITION 1.1.21. A BRT signature is

- i) a finite list of one or more distinct core BRT terms; or
- ii) a finite list of one or more distinct core BRT terms, followed by the symbol \subseteq .

DEFINITION 1.1.22. The entries of a BRT signature are just its core BRT terms.

Let α be a BRT fragment. I.e., let $\alpha = "[Environment] \text{ in } \sigma \text{ on } [Setting]"$ be a BRT fragment, where σ is a BRT signature.

DEFINITION 1.1.23. The signature of α is σ . The α terms are defined by

- i) every entry of σ is an α term;
- ii) \cup, \emptyset are α terms;
- iii) if s, t are α terms then $(s \cup t), (s \cap t), (s \setminus t)$ are α terms.

The α terms are to be distinguished from the entries of σ , since we are closing the entries of σ under Boolean operations.

DEFINITION 1.1.24. The α equations are the equations between α terms. The α inequations are the inequations (\neq) between α terms. The α inclusions are the inclusions between α terms.

DEFINITION 1.1.25. The α formulas are inductively defined by

- i) every α equation is an α formula;
 ii) if φ, ψ are α formulas, then $(\neg\varphi), (\varphi \vee \psi), (\varphi \wedge \psi), (\varphi \rightarrow \psi), (\varphi \leftrightarrow \psi)$ are α formulas.

DEFINITION 1.1.26. The α basics are the α equations if the environment of α is EBRT; the α inequations if the environment of α is IBRT; the α formulas if the environment of α is PBRT.

Suppose first that the signature σ of α does not end with $\underline{\subseteq}$. Let the BRT setting of α be (V, K) .

DEFINITION 1.1.27. An α assignment is an assignment of an element of V to each function variable appearing in σ , and an element of K to each set variable appearing in σ .

DEFINITION 1.1.28. The α assertions are assertions of the form

$$(\forall g_1, \dots, g_n \in V) (\exists B_1, \dots, B_m \in K) (\varphi)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and φ is an α basic.

Now assume that σ ends with $\underline{\subseteq}$.

DEFINITION 1.1.29. An α assignment is an assignment of an element of V to each function variable appearing in σ , and an element of K to each set variable appearing in σ , where if A_i, A_j appear in σ and $1 \leq i \leq j$, then the set assigned to A_i is included in the set assigned to A_j .

DEFINITION 1.1.30. The α assertions are assertions of the form

$$(\forall g_1, \dots, g_n \in V) (\exists B_1 \subseteq \dots \subseteq B_m \in K) (\varphi)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and φ is an α basic.

Thus if the environment of α is EBRT, then the α assertions are based on α equations φ . If the environment of α is IBRT, then the α assertions are based on α inequations φ .

If the environment of α is PBRT, then the α assertions are based on α formulas φ . These hold regardless of whether the signature of α ends with \subseteq .

DEFINITION 1.1.31. We say that an α formula is α valid if and only if it holds for all α assignments.

DEFINITION 1.1.32. Let φ, ψ be α formulas. We say that φ, ψ are α equivalent if and only if $\varphi \leftrightarrow \psi$ is α valid. This definition is extended to sets of α formulas in the obvious way.

This concludes the definition of BRT fragments, and their assertions.

The above treatment of BRT fragments, $\alpha =$

[Environment] in [Signature] on [Setting]

fully explains the titles of the Classification sections 2.4 - 2.7.

DEFINITION 1.1.33. The standard BRT signatures have the form

$$A_1, \dots, A_n, f_1A_1, \dots, f_1A_n, \dots, f_mA_1, \dots, f_mA_n \\ A_1, \dots, A_n, f_1A_1, \dots, f_1A_n, \dots, f_mA_1, \dots, f_mA_n, \subseteq$$

and are referred to as

$$m \text{ functions and } n \text{ sets.} \\ m \text{ functions and } n \text{ sets}/\subseteq.$$

where $n, m \geq 1$. A flat BRT signature is a BRT signature where every entry is either some A_i , or some f_iA_j , or some f_iU .

DEFINITION 1.1.34. A standard BRT fragment is a BRT fragment whose environment is EBRT or IBRT, and whose signature is a standard BRT signature. A flat BRT fragment is a BRT fragment, with environment EBRT or IBRT, whose signature is flat.

The BRT fragments considered in sections 2.2, 2.4-2.7, and Chapter 3, are all standard. In section 2.3, we work with the flat signature A, fA, fU . In Chapter 3, we are successful in analyzing a small part of the standard BRT fragment (see Definition 2.1)

EBRT in $A, B, C, fA, fB, fC, gA, gB, gC$ on (ELG, INF) .

For example, in this book we do not consider such interesting BRT signatures as

$$\begin{aligned} & A, fA, ffA. \\ & A, fA, f(U \setminus A). \\ & A, B, fA, fB, f(A \cup B), \subseteq. \end{aligned}$$

none of which are flat.

Let α be a standard BRT fragment with m functions and n sets, whose signature does not end with \subseteq . Then the number of entries of the signature is $n(m+1)$. So the number of α terms is $2^{2^{n(m+1)}}$ up to Boolean identities. Therefore the number of α basics is also $2^{2^{n(m+1)}}$ up to formal Boolean equivalence. This is also the number of α assertions up to formal Boolean equivalence.

The number of α assertions, up to formal Boolean equivalence, grows very rapidly. For 1 function and 1 set, we have $2^{2^2} = 16$. For 1 function and 2 sets, we have $2^{2^4} = 2^{16} = 65,536$. For 1 function and 3 sets, we have $2^{2^6} = 2^{64}$. For 2 functions and 2 sets, we have $2^{2^6} = 2^{64}$. For the second, third, and fourth of these cases, we do not know if the α assertions on the basic BRT settings considered here include assertions independent of ZFC. We believe that they do not.

The number of α assertions grows less rapidly, up to BRT equivalence, if the signature ends with \subseteq . This reduction of complexity allows us to work successfully with EBRT in A, B, fA, fB, \subseteq on various basic settings, in Chapter 2.

For standard BRT fragments with 2 functions and 3 sets, without \subseteq in the signature, we have $2^{2^9} = 2^{512}$ assertions. The so called Principal Exotic Case lives in EBRT in the standard signature with 2 functions and 3 sets, on the BRT setting (ELG, INF) . The Principal Exotic Case is Proposition A from Appendix A, and is the focus of Chapters 4 and 5 where it is shown to be independent of ZFC (assuming SMAH = ZFC augmented with the existence of strongly Mahlo cardinals of each finite order, is consistent).

The Principal Exotic Case lies formally in the standard BRT fragment

EBRT in $A, B, C, fA, fB, fC, gA, gB, gC$ on (ELG, INF) .

In fact, the Principal Exotic Case lives in the considerably reduced flat BRT fragment

EBRT in A, C, fA, fB, gB, gC on (ELG, INF) .

In fact, we can strengthen the Principal Exotic Case with $A \subseteq B \subseteq C$, which now lives in the further reduced flat BRT fragment

EBRT in $A, C, fA, fB, gB, gC, \subseteq$ on (ELG, INF) .

In Chapters 4 and 5, we show that both of these statements are provable using large cardinals, but not in ZFC (assuming ZFC is consistent).

It is important to have a useful format for presenting BRT assertions. For the purposes of Chapter 2, this amounts to creating a useful format for presenting BRT equations. The most useful format is a set of pre elementary inclusions, or a set of elementary inclusions, defined below.

Let α be a flat BRT fragment, with signature σ .

DEFINITION 1.1.35. The α pre elementary inclusions are of the form

- i) $t_1 \cap \dots \cap t_n = \emptyset$, where $n \geq 1$, t_1, \dots, t_n are the entries of σ , in order of their appearance in σ ;
- ii) $t_1 \cup \dots \cup t_n = U$, where $n \geq 1$, t_1, \dots, t_n are the entries of σ , in order of their appearance in σ ;
- iii) $r_1 \cap \dots \cap r_p \subseteq s_1 \cup \dots \cup s_q$, where $p, q \geq 1$, and $r_1, \dots, r_p, s_1, \dots, s_q$ are a listing of all of the entries of σ without repetition, and r_1, \dots, r_p and s_1, \dots, s_q are both in order of their appearance in σ .

Note that if there are n entries of σ , then there are 2^n α pre elementary inclusions.

DEFINITION 1.1.36. Suppose σ does not end with \subseteq . The α elementary inclusions are obtained from the σ pre elementary inclusions in the following way. If fU and some fA_i appears in an intersection, then remove fU there. If fU appears in a union, then remove all fA_i there. In order to

be an elementary inclusion, we require that for every fA on the left, fU must not be on the right.

Note that if fU is not an entry of the signature of α , then the elementary inclusions are just the pre elementary inclusions.

Now suppose the signature of α ends with \subseteq .

DEFINITION 1.1.37. Suppose σ ends with \subseteq . The α elementary inclusions are obtained from the α pre elementary inclusions in the following way. For any A appearing in an intersection, retain only the A_i where i is least. For any A appearing in a union, retain only the A_i where i is greatest. For any f appearing in an intersection, retain only the fA_i where i is least (if only fU appears, then retain fU). For any f appearing in a union, retain the fA_i where i is greatest (if fU appears, then retain only fU). In order to be an elementary inclusion, we require that for every fA_i on the left, fU must not be on the right, and every fA_j , $j \geq i$, must not be on the right.

DEFINITION 1.1.38. An α format is a set of α elementary inclusions.

In case σ does not end with \subseteq , our α formats take advantage of the fact that $fA_i \subseteq fU$. In case σ ends with \subseteq , our α formats take advantage of the fact that $A_i \subseteq A_j$ and $fA_i \subseteq fA_j \subseteq fU$, for $i < j$.

We need to verify that our reduction to α formats is valid; i.e., covers what we want. This amounts to verifying that every α equation is α equivalent to an α format. In fact, we show that every set of α inclusions is α equivalent to an α format.

THEOREM 1.1.1. Let α be a flat BRT fragment. Every α inclusion is α equivalent to an α format. Every set of α inclusions is α equivalent to an α format. Every α format is α equivalent to an α inclusion, and an α equation.

Proof: We first assume that the signature of α does not end with \subseteq .

For the first claim, let $s \subseteq t$ be an α inclusion. Using standard Boolean algebra, write s as a union of intersections of entries and complements of entries of the

signature σ . Write t as an intersection of unions of entries and complements of entries of σ . Here the complements are taken with respect to the universal set U . We allow the degenerate case where s is \emptyset, U , and t is \emptyset, U . Of course, intersections and unions of cardinality 1 are also allowed.

We then obtain a set of inclusions $s' \subseteq t'$, where the s' are intersections of entries and complements of entries from σ , and the t' are unions of entries and complements of entries of σ . Again, we allow the degenerate case of $s', t' = \emptyset, U$. We can remove all such degenerate cases except $U \subseteq \emptyset$.

We can now arrange for each of these inclusions to be of the forms

$$\begin{aligned} \pm s_1 \cap \dots \cap \pm s_n &\subseteq \pm t_1 \cup \dots \cup \pm t_m. \\ U &\subseteq \pm t_1 \cup \dots \cup \pm t_m. \\ \pm s_1 \cap \dots \cap \pm s_n &\subseteq \emptyset. \\ U &\subseteq \emptyset. \end{aligned}$$

And then of the forms

$$\begin{aligned} \pm s_1 \cap \dots \cap \pm s_n &\subseteq \pm t_1 \cup \dots \cup \pm t_m. \\ \pm t_1 \cup \dots \cup \pm t_m &= U \\ \pm s_1 \cap \dots \cap \pm s_n &= \emptyset. \\ U &= \emptyset. \end{aligned}$$

Here $n, m \geq 1$, and the s 's and t 's are entries in σ . We write $+t$ for t and $-t$ for $U \setminus t$. We must allow for the possibility that there are no inclusions. This corresponds to the case where we have only $U = U$.

We can also require that in each of these clauses, each s_i can appear only once, each $-s_i$ can appear only once, and we cannot have s_i and $-s_i$ appear. This is because of the Boolean equivalence

$$X \cap Y \subseteq Z \cup -Y \leftrightarrow X \cap Y \subseteq Z.$$

We now replace each of the above five forms with an equivalent set of inclusions in which all entries of σ appear (or their complement). Thus suppose

$$\pm s_1 \cap \dots \cap \pm s_n \subseteq \pm t_1 \cup \dots \cup \pm t_m$$

is missing $\pm r_1, \dots, \pm r_k$. Then replace it with the set of all

$$\pm s_1 \cap \dots \cap \pm s_n \subseteq \pm t_1 \cup \dots \cup \pm t_m \cup \beta_1 \cup \dots \cup \beta_k$$

where each β_i is r_i or $-r_i$.

Suppose

$$\pm t_1 \cup \dots \cup \pm t_m = U$$

is missing entries $\pm r_1, \dots, \pm r_k$. Then replace it with the set of all

$$\pm t_1 \cup \dots \cup \pm t_m \cup \beta_1 \cup \dots \cup \beta_k = U$$

where each β_i is r_i or $-r_i$.

Suppose

$$\pm s_1 \cap \dots \cap \pm s_n = \emptyset.$$

is missing entries $\pm r_1, \dots, \pm r_k$. Then replace it with the set of all

$$\pm s_1 \cap \dots \cap \pm s_n \subseteq \beta_1 \cup \dots \cup \beta_k$$

where each β_i is r_i or $-r_i$.

Replace $U = \emptyset$ with

$$\beta_1 \cup \dots \cup \beta_k = U$$

where each $\beta_i = r_i$ or $-r_i$ and r_1, \dots, r_k is a list without repetition of all entries of σ .

We now have a set of what would be α pre elementary inclusions except for the fact that complements are present. However, we can eliminate the complements by shifting from one side to the other according to the following Boolean equivalences.

$$\begin{aligned} X \subseteq Y \cup U \setminus Z &\leftrightarrow X \cap Z \subseteq Y. \\ X \cap U \setminus Y \subseteq Z &\leftrightarrow X \subseteq Y \cup Z. \\ X \subseteq U \setminus Z &\leftrightarrow X \cap Z = \emptyset. \\ U \setminus Y \subseteq Z &\leftrightarrow Y \cup Z = U. \end{aligned}$$

Thus we are left with a set (possibly empty) of α pre elementary inclusions.

Recall the process of converting any α pre elementary inclusion to an α elementary inclusion. The given α pre elementary inclusion is obviously α equivalent to the resulting α elementary inclusion. Thus we are left with an α format.

This establishes the first claim. The second claim follows immediately from the first claim since a finite set of α inclusions can be written as a single α inclusion.

For the final claim, let $\{s_1 \subseteq t_1, \dots, s_n \subseteq t_n\}$ be an α format, $n \geq 0$. If $n = 0$, then take $A \subseteq A$, $A = A$, where A is an entry of σ . Suppose $n > 0$. Then use the Boolean equivalence

$$\begin{aligned} s_1 \subseteq t_1 \wedge \dots \wedge s_n \subseteq t_n &\leftrightarrow \\ s_1 \setminus t_1 \cup \dots \cup s_n \setminus t_n = \emptyset &\leftrightarrow \\ s_1 \setminus t_1 \cup \dots \cup s_n \setminus t_n \subseteq \emptyset. & \end{aligned}$$

We now assume that the signature of α does end with \subseteq . Let α be $\alpha' \subseteq$. For the first claim, let $s \subseteq t$ be an α inclusion. As before, we obtain an equivalent set of pre elementary inclusions for α' . At this point, we perform the reductions that create an equivalent set of pre elementary inclusions for α . We then proceed as above to create an equivalent set of elementary inclusions for α .

The second and third claims are proved as before. QED

We will use Theorem 1.1.1 as follows. Let α be a flat BRT fragment with signature σ and BRT setting (V, K) .

Suppose the environment of α is EBRT, and α does not end with \subseteq . By Theorem 1.1.1, the α assertions can be put into the form

$$(\forall g_1, \dots, g_n \in V) (\exists B_1, \dots, B_m \in K) (S)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and S is an α format, interpreted conjunctively.

Suppose the environment of α is EBRT, and σ ends with \subseteq . By Theorem 1.1.1, the α assertions can be put into the form

$$(\forall g_1, \dots, g_n \in V) (\exists B_1 \subseteq \dots \subseteq B_m \in K) (S)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and S is an α format, interpreted conjunctively.

Suppose the environment of α is IBRT, and σ does not end with \subseteq . To avoid considering the very awkward negated formats, we work with the dual. Thus the inequation becomes an equation, so that we can apply Theorem 1.1.1. By Theorem 1.1.1, the α assertions can be put into the form

$$\neg (\exists g_1, \dots, g_n \in V) (\forall B_1, \dots, B_m \in K) (S)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and S is an α format, interpreted conjunctively.

Suppose the environment of α is IBRT, and σ does end with \subseteq . By Theorem 1.1.1, the α assertions can be put into the form

$$\neg (\exists g_1, \dots, g_n \in V) (\forall B_1 \subseteq \dots \subseteq B_m \in K) (S)$$

where $n, m \geq 1$, B_1, \dots, B_m are the BRT set variables mentioned in σ with strictly increasing subscripts, g_1, \dots, g_n are the BRT function variables mentioned in σ with strictly increasing subscripts, and S is an α format, interpreted conjunctively.

As indicated above, Theorem 1.1.1 tells us that we need only work with

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

where the g 's and B 's are as indicated earlier, and S is an α format. It will be seen to be very convenient to drop the negation signs in front of the last two of the above.

DEFINITION 1.1.39. Let α be a flat BRT fragment. The α statements (rather than the α assertions) are statements of form 1) above if the environment of α is EBRT and the signature of α does not end with \subseteq ; 2) above if the environment of α is EBRT and the signature of α ends with \subseteq ; 3) above if the environment of α is IBRT and the signature of α does not end with \subseteq ; 4) above if the environment of α is IBRT and the signature of α ends with \subseteq .

DEFINITION 1.1.40. Let α be a flat BRT fragment. An α format S is said to be correct if and only if the α statement using S is true; incorrect otherwise.

Informally speaking, a classification of a BRT fragment α amounts to a determination of all α correct α formats.

As discussed earlier, the number of pre elementary inclusions in the standard signature

$$A_1, \dots, A_n, f_1A_1, \dots, f_1A_n, \dots, f_mA_1, \dots, f_mA_n$$

with m functions and n sets is $2^{n(m+1)}$, and the number of formats is therefore $2^{2^{n(m+1)}}$.

THEOREM 1.1.2. The number of elementary inclusions in $A_1, \dots, A_n, f_1A_1, \dots, f_1A_n, \dots, f_mA_1, \dots, f_mA_n, \subseteq$ is $(n+1)^{m+1}$. Therefore the number of formats (or statements) is $2^{(n+1)^{m+1}}$. In the case of $A, B, C, fA, fB, fC, gA, gB, gC, \subseteq$, we have 64 and 2^{64} . In the case of $A, C, fA, fB, gB, gC, \subseteq$, we have 27 and 2^{27} . In the case of A, B, fA, fB, \subseteq , we have 9 and 2^9 .

Proof: Let us first focus on the pattern of A 's in elementary inclusions. Recall that the elementary inclusions are the immediate simplifications of the pre elementary inclusions, using $A_1 \subseteq \dots \subseteq A_n$.

- i. A_i on left, A_{i-1} on right.
- ii. A_1 on left, no A_j on right.
- iii. No A_i on left, A_n on right.

There are $n+1$ among i-iii. The same count holds for the other m groups. So we obtain a total of $(n+1)^{m+1}$ elementary inclusions. QED

For PBRT in σ on (V,K) , where σ is based on m functions and n sets, we cannot specify the assertions by a single format. Instead, what is relevant is the number of all α formulas up to propositional and Boolean equivalence. The number of α equations, up to Boolean equivalence, is $2^{2^n(m+1)}$, and so the number of α formulas up to propositional and Boolean equivalence is $2^{2^{2^n(m+1)}}$. This quantity is quite frightening. Even in one function and one set, this is $2^{2^{2^2}} = 2^{65,536}$. For one function and two sets, this is $2^{2^{2^4}} = 2^{2^{65,636}}$. These numbers do not address, say, two functions and three sets. We do not tackle PBRT in this book.

In Chapter 2, we focus on the five basic BRT settings, (SD, INF) , $(ELG \cap SD, INF)$, (ELG, INF) , $(EVSD, INF)$, and (MF, INF) .

In section 2.2, we classify EBRT/IBRT in A, fA on the five basic BRT settings, where the number of assertions is $2^{2^2} = 16$. This is of course completely manageable, but still turns out to be substantial. Already, the significant Thin Set Theorem and Complementation Theorem appear among the 16.

In section 2.3, we classify EBRT/IBRT in A, fA, fU on the five basic BRT settings, where the number of statements is $2^{2^3} = 256$, with considerable duplication due to equivalence on all BRT settings. This is still manageable.

For EBRT/IBRT in A, B, fA, fB , the number of statements is $2^{2^4} = 2^{16} = 65,536$. This is rather daunting, but within manageability with a few years of effort. This optimism is based on the expectation that there will be a large proportion of trivial cases, and lots of relations between cases. This has been the experience with sections 2.4 and 2.5.

In sections 2.4 and 2.5, we classify EBRT in $A, B, fA, fB, \underline{C}$ on the five basic settings excluding (MF, INF) , where the number of statements is $2^9 = 512$ (according to Theorem 1.1.2). As can be seen from sections 2.6 and 2.7, we can go much further in the fifth basic BRT setting, (MF, INF) , as well as in IBRT on all five basic settings.

In sections 2.2 and 2.3, we make a brute force enumeration of cases. However, in sections 2.4 - 2.7, we prefer to use a treelike methodology. This treelike methodology is presented in section 2.1, where we also develop the relevant theory.

We see that all of the BRT statements that arise from the EBRT classifications in Chapter 2 are decided in RCA_0 , and all of the BRT statements that arise from the IBRT classifications in Chapter 2 are decided in ACA' .

ZFC incompleteness arises somewhat later in the development of BRT, with EBRT in $A, B, C, fA, fB, fC, gA, gB, gC$ on the BRT setting (ELG, INF) . The Principal Exotic Case, also known as Proposition A in this book, lies within this BRT fragment (see Appendix A). Here we have $2^{2^9} = 2^{512}$ statements. This is entirely unmanageable. It would take several major new ideas to make this manageable in any sense of the word. The same is true even for $A, B, C, fA, fB, fC, gA, gB, gC, \subseteq$, since by Theorem 1.1.2, this involves 2^{64} statements. There is a lot of simplification coming from \subseteq , but there does not seem to be nearly enough for manageability.

However, the Principal Exotic Case lies within the much smaller fragment EBRT in $A, C, fA, fB, gB, gC, \subseteq$ on (ELG, INF) . We expect to get enough substantive simplification from \subseteq to make $A, C, fA, fB, gB, gC, \subseteq$ as a manageable decade long project. According to Theorem 1.1.2, the relevant count is 2^{27} before substantive simplifications.

In section 3, we give a classification for a very restricted subclass of the statements for EBRT in $A, B, C, fA, fB, fC, gA, gB, gC$ on the BRT setting (ELG, INF) . The Principal Exotic Case lies within this very restricted subclass with $3^8 = 6561$ statements.

The Principal Exotic Case is shown in Chapters 4, 5 to be provable using strongly Mahlo cardinals of all finite orders, yet not provable in ZFC (assuming ZFC is consistent).

We have given only an informal account of what we mean by a classification for a BRT fragment. We now seek to be more formal.

DEFINITION 1.1.41. Let α be a BRT fragment. A tabular classification for α is a table of the correct α formats.

However, this definition does not take into account the background theory needed to document the table, which is of importance for BRT.

Let T be a formal system with an adequate definition of the BRT fragment α .

DEFINITION 1.1.42. We say that an α format S is α, T correct if and only if the α statement using S is provable in T . We say that S is α, T incorrect if and only if the α statement using S is refutable in T .

DEFINITION 1.1.43. We say that α is T secure if and only if every α format is α, T correct or α, T incorrect.

DEFINITION 1.1.44. A tabular α, T classification consists of a table of all α formats, together with a proof or refutation of each of the corresponding α statements, within T . This is a rather direct demonstration that α is T secure.

In sections 2.2, 2.3, we provide what amounts to a tabular α, T classification for some simple BRT fragments α , where T is RCA_0 or a weak extension of RCA_0 .

In Chapter 3, we provide what amounts to a tabular $\alpha, SMAH^+$ classification for a very limited subclass of the EBRT formats α in $A, B, C, fA, fB, fC, gA, gB, gC$ on (ELG, INF) . (For $SMAH^+$, see Appendix A).

But for some α , it is not reasonable to present such large tables. How do we show that α is T secure? What then do we mean by a classification of the α statements in T ?

In sections 2.4 - 2.7, we do not present tables, but instead use a treelike methodology. In section 2.1, we develop the theory of this methodology, showing how the analyses in sections 2.4 - 2.7 demonstrate that α is T secure, for various α, T .

To give a classification of the α statements in T , it suffices to give a listing of the maximally α, T correct α formats; i.e., the α, T correct α formats that are not properly included in any other α, T correct format.

In section 2.1, we define the T classifications, TREE, for α . We prove that there is a T classification for α if and only if α is T secure. We also give an algorithm for generating the maximally α, T correct α formats from TREE. We show that the number of maximally α, T correct α formats is at most the number of vertices in TREE.

In sections 2.4 - 2.7, T classifications for the relevant α are actually given in a style prescribed in section 2.1, where $T = RCA_0$ and ACA' .

The classifications given in Chapters 2 and 3 are rather limited in scope. For instance, we conjecture that

- i. EBRT in A, B, fA, fB, fN on any of the five basic BRT settings is RCA_0 secure.
- ii. IBRT in A, B, fA, fB, fN on any of the five basic BRT settings is ACA' secure.
- iii. EBRT/IBRT in $A, B, C, fA, fB, fC, gA, gB, gC, fN, gN$ on any of the five basic BRT settings is $SMAH^+$ secure.

These conjectures are wide open. In fact, we have not even established any of i-iii for A, B, fA, fB . We have established i-iii for $A, B, fA, fB, \underline{C}$.