A Topological Framework for Modelling Belief Revision

by

Lindsey Craig Jeftha

Dissertation presented for the degree of Doctor of Philosophy

at

Stellenbosch University

Department of Mathematical Sciences

Faculty of Science

Private Bag X1, Matieland, 7602, South Africa

Promoter: Professor Ingrid Rewitzky

December 2010

Declaration

By submitting this dissertation electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the owner of the copyright thereof (unless to the extent explicity otherwise stated) and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Signed: ______________________________
L.C. Jeftha

Date: ______________________________

Copyright © 2010 Stellenbosch University

All rights reserved

Abstract

Classical formulations model belief revision as a deterministic process. Under certain circumstances, the process may have more than one outcome, which suggests that belief revision is non-deterministic instead. Representations exist that model belief revision in either format, and for both formats there are axiom schemes that determine whether the representation is in fact a belief revision process.

Although the axiom scheme for the non-deterministic case generalises that of the deterministic case, both schemes entail that all of the beliefs held by an agent are affected by new information, which is perhaps unintuitive. Rather, one may consider that belief revision should be local, with beliefs only affected if the new information is pertinent to them. We approach the problem of belief revision from the standpoint that it is local and non-deterministic, and the purpose and contribution of this dissertation is the development of a topological framework with which to model belief revision in this manner.

Opsomming

Geloofshersiening word gewoonlik as 'n deterministiese proses voorgestel. Meer as een uitkoms mag bestaan vir verskeie omstandighede, wat aandui dat die proses liewer nie-deterministies van aard is. Beide die gevalle word deur aksiomaskemas gereguleer, en die aksiomas vir die nie-deterministiese geval veralgemeen dié van die deterministiese geval.

Albei aksiomaskemas stipuleer, miskien onintuïtief, dat alle gelowe van 'n agent deur die nuwe informasie geaffekteer word. 'n Beter metode is dat net daardie gelowe waarvoor die nuwe informasie toepaslik is geaffekteer word. Ons benader die probleem van geloofshersiening uit die standpunt dat dit lokaal en nie-deterministies is, en die doel en bydrae van hierdie proefskrif is dus die ontwikkeling van 'n topologiese raamwerk waarmee ons geloofshersiening op hierdie manier kan voorstel.

Acknowledgements

The preparation of this dissertation has been a wonderful adventure, and a real journey of discovery. Sometimes gruelling, mostly challenging, but always interesting.

Without the patient support and encouragement of my mother and father – not to mention the laundry and many wonderful free meals I got to enjoy – I would not have made it through the experience in one piece.

Thanks are also due to my dear aunt, Lucy, whose belief and confidence in my ability to reach this point were astounding, and a source of encouragement during many of the darker moments that inevitably accompany a journey such as this.

Thank you also to both my grandmothers, to my paternal grandmother for her kindness, patience and faith, and for reading me all those stories about aeroplanes and pilots when I was younger, and to my late maternal grandmother for her example, high standards and insistence on the truth that have always been a beacon along life's way.

Thank you to my piano teacher, Professor Graham Fitch. Mentioning your name in a text such as this does seem incongruous, however my piano lessons with you were always very stimulating and provided a welcome and essential diversion, even sanctuary, from the "alternate reality" that is mathematics. Thank you for always providing a warm, welcoming and stimulating learning environment, and most especially for always listening to and taking the time to answer my questions. I will certainly remember these lessons and your insightful teaching with much gratitude.

I also owe a debt of gratitude to Gabriël, Andrew and Ralf. To Gabriël, with whom I have been lucky enough to share many fruitful discussions about subjects as far-ranging as mathematics, religion and mountain-climbing, thank you for patiently enduring my endless questions! I am truly fortunate to have a friend like you. To Andrew, thank you for keeping me on my toes, and to Ralf, thanks for being an all round great and very funny friend, and to all of you for being such great fun to hang out with.

Thank you to my supervisor, Professor Ingrid Rewitzky, who helped with the review and correction of my work and also very kindly arranged for partial financial support from the NRF, whose contribution I wish gratefully to acknowledge.

Finally, and most importantly, thanks to my best friend, for once again being the duct tape, paper clips and chewing gum that have kept my ship together for the past 9 months. I am more grateful than I can say, and in a "refutable" sense, there is not enough paper in the world for me to thank You properly for all of Your help.

Contents

Chapter 1

A Survey of Belief Revision

Over time, an agent will accumulate a store of beliefs about its environment. When it encounters new information about this environment, we can certainly expect it to change its beliefs to be consistent with the new information in case some of its beliefs are invalidated. The mechanism of this change is not straightforward, as there are many technical difficulties to overcome. The problem has attracted much attention as a result, culminating in the publication of the seminal works [1] and [30] ([31]), with a very large body of subsequent research that covers a wide variety of innovative approaches to the problem.

Within this body of research, two kinds of belief change are generally distinguished [6, 26, 30], *viz.* belief update and belief revision. Many researchers feel that the distinction between update and revision is not clear (compare [32, 6]), so for our work we shall concentrate on belief revision. In this chapter, we present an overview of three different approaches to belief revision, after which we are able to describe what it is that we aim to accomplish through our work. We proceed as follows.

Chapter Guide:

Section 1.1: Belief, Knowledge and Acceptance. In this section, we summarise some contemporary views on what it means to believe, know or accept something. It is important to understand these concepts because they help to shape our expectation of what should happen when beliefs change to accommodate new information.

Section 1.2: Different Approaches to Belief Revision. Belief revision can be viewed as a transition between different states of affairs. In this section we describe three approaches to modelling these transitions. The first approach models belief revision as a functional operation and is set out in [1]. The second approach, proposed in [11], refines the approach of [1] to account for relevance. In the third approach, described in [36], belief revision is modelled as a relational operation.

Section 1.3: The Aims of Our Work. We now use the ideas of the preceding sections to formulate our thesis. We describe what it is that we intend to accomplish and contribute, and present an overview of the remainder of our work.

1.1 Belief, Knowledge and Acceptance

To believe something entails that there is an agent that holds the given belief, and that there is something that this agent believes, the content of the belief. Traditionally, this content is a proposition expressed in some chosen language [45]. Belief then refers to the representation held by an agent of the truth value of a proposition [24]. This truth value is independent of its representation, for a proposition can be false even though an agent believes it to be true. The beliefs held by an agent are then just the propositions that it considers to be true.

We may treat this set of beliefs as a state of affairs or as a possible world, as conceived of by the agent. Many contemporary philosophers then characterise belief as the implicitly modal concept of a propositional attitude [45]. A propositional attitude is an opinion about or a disposition towards a proposition or the state of affairs in which that proposition is true. It is often expressed as "$A(x, \phi)$", where A is the propositional attitude, x is an individual and ϕ is a sentence that expresses the proposition. For example, "James is hopeful that there is enough food for everyone" has "James" as x, "is hopeful that" as A and "there is enough food for everyone" as ϕ.

The subject of knowledge and what it means to say that "x knows that ϕ" is studied in a broad-ranging field of philosophy called epistemology. It encompasses questions such as "What is knowledge?" and "How do people acquire knowledge?", and probes distinctions between "knowing *how*" and "knowing *that*". We shall concern ourselves only with "knowing that".

The nature and definition of knowledge are elusive, even intractable (compare [51]), and debates on the subject are beset by controversy. Much of the controversy stems from the formulation of knowledge as justified true belief, where "x knows that ϕ" if and only if ϕ is the case, x believes that ϕ and x is justified in believing that ϕ [51]. This formulation was criticised famously in [21] through a suite of counterexamples known as the "Gettier Problems", which illustrate situations where a belief may be justified and true yet not count as knowledge.

In a typical counterexample, a belief held by an agent coincides with the true situation by chance. This justified, true belief is then falsely classified as knowledge. For example, based on [51], suppose that Jack believes the false proposition p, "Mary owns a Ford". Suppose further that Jack infers proposition q, "either Mary owns a Ford or John is in Barcelona". Since p entails q, Jack is justified in believing q. If John just happened to be in Barcelona, q would be true, and Jack would hold a justified, true belief. However, Jack has no evidence for the whereabouts of John, so his justified true belief cannot be regarded as knowledge.

Since publication of the Gettier problems, much debate around the nature of knowledge has focussed on finding a suitable clause to add to the original formulation of knowledge as justified true belief that will "de-gettierise" it. Nevertheless, the nature of knowledge remains at large. In particular, with regard to the debate about the nature of knowledge, [51] expresses the opinion that

> "One way to respond to the intractability of the debate is to acknowledge that there simply is not *one* concept of knowledge for which there is an analysis that has any chance of meeting with broad assent. Rather, we might conclude that, when we use the word 'knowledge', we have sometimes one concept and at other times another concept in mind."

However we wish to characterise knowledge, it seems reasonable that what is false cannot be known (compare [51]) – a statement such as "Roald Dahl wrote 'Pride and Prejudice' " is false, so it is not something that one could claim to know. Furthermore, it is nearly universally accepted that knowledge implies belief, and it is taken as a contradiction to claim to know something without believing it (compare p118 in [43] and also [53]).

From [45], acceptance is taken to be under the voluntary control of the subject more so than belief and is more directly linked to a particular action and context. One can believe something without accepting it, and also accept something without believing it. For example, given evidence that supports a theory but is not known to be completely decisive, a scientist may choose to accept or disregard the theory, and furthermore may do so without needing to believe it (a similar version of this example can be found in [18] ([19])).

In contrast, suppose that Bill needs to climb a ladder to clean the gutters around his house. He may genuinely believe that the ladder is stable and can support his weight, but out of concerns for safety he does not accept this to be so until he has properly checked the ladder. Thus Bill believes that the ladder is safe without accepting it to be so (a similar example to this can be found in [45]).

To accept something seems to involve a choice to cease enquiry and treat the matter as settled at least until new evidence becomes available. Literature on the subject is generally not clear how acceptance relates to knowledge and belief. For example, can one know something without accepting it? If knowledge is always believed, then because we can believe something without accepting it, arguably we can also know something without accepting it. To illustrate, one could know that air travel is safe without accepting it because accidents can happen. Yet, one cannot know something without its being the case, so to know something seems to involve our acceptance of it also.

From our brief overview of belief, knowledge and acceptance, many of the difficulties we encounter are yet mired in active though apparently intractable debate. In the study of belief change however, several advances have been made through the innovative approaches taken by different researchers, and in the next section we shall examine some of these approaches in more detail.

1.2 Different Approaches to Belief Revision

Generally, two kinds of belief change are distinguished [6, 26], *viz.* belief update, where the beliefs of an agent are modified in response to changes in its environment, and belief revision, where the beliefs of an agent are modified when it receives new information about its unchanged environment.

For both kinds of belief change, there is the pervasive intuition that a body of beliefs should undergo minimal changes to accommodate new information (compare [42] however, where even this widely accepted notion is questioned). In [1], the authors propose a suite of postulates, known as the "AGM Postulates", which formalise this notion of minimal change for the case of belief revision and which have become the dominating paradigm for reasoning about belief revision. The authors of [30] ([31]) have similarly proposed a collection of postulates, known as the "KM Postulates", that regulate the operation of belief update.

Many researchers feel that the distinction between update and revision is not clear (compare [32]) and that belief update contains elements of belief revision as well (compare [6]). On the one hand, differences between the operations make them largely incompatible, so that one cannot be treated as a special case of the other [26]. On the other hand, in [6] the author provides an example to show that belief update can also contain elements of belief revision, and then exhibits semantics for a model that seeks to unify the two types of belief change.

In view of the similarities between belief update and belief revision, we shall concentrate on belief revision only. Belief revision is sometimes further characterised in terms of the related operations of contraction and expansion, whereby a piece of information is removed from (resp. added to) a body of beliefs. In general, revision is regarded as combining these operations (compare [35]), so in our work we will not additionally study expansion and contraction.

A belief revision operation may be thought of as a transition from one state of affairs to another, and in the rest of this section we examine three approaches to modelling these transitions. The first approach models belief revision as a functional operation and is described in [1]. The second approach, proposed in [11], refines the approach of [1] to account for relevance. The third approach models belief revision as a relational operation, and is described in [36].

1.2.1 The Functional Approach

Under the functional approach, a belief revision operation $*$ is a function that takes a deductively closed set K of beliefs and formula ϕ to the deductively closed set $K * \phi$ of beliefs. The authors of [1] propose that the properties of $*$ and the new beliefs $K * \phi$ are governed by the postulates **A1 – A8**, the AGM Postulates, listed below.

To formulate these postulates, we fix a propositional though not necessarily finitary language L. The well-formed formulae of L are constructed by application of the binary connectives $\wedge$ (AND) and $\vee$ (OR) and the unary connective $\neg$ (NOT) in the standard manner, and we collect these formulae in the set Φ_L. We fix a set-to-set function C_n that maps a set Φ of well-formed formulae of L to the set Ψ of formulae that can be deduced from Φ. We denote a contradiction in Φ_L by $\perp$, and the entailment relation by $\vdash_L$. A theorem of the language is then a formula ϕ such that $\vdash_L \phi$. We take $C_n(\{\perp\}) = \Phi_L$. A set K of formulae is then consistent if $C_n(K) \neq C_n(\{\perp\}) = \Phi_L$ (compare Definition 3.9(ii) and p115 in [12]), otherwise it is inconsistent. A belief set K is a set of formulae expressed in L that is closed under deduction, *i.e.* $K = C_n(K)$.

The AGM postulates may then be formulated as follows (compare [1] and [20]):

A1: $K * \phi$ is a belief set

A2: $\phi \in K * \phi$

A3: $K * \phi \subseteq C_n(K \cup \{\phi\})$

A4: If $\neg\phi \notin K$ then $C_n(K \cup \{\phi\}) \subseteq K * \phi$

A5: $K * \phi = C_n(\{\perp\})$ if and only if $\vdash_L \neg\phi$

A6: If $\vdash_L \phi \leftrightarrow \psi$ then $K * \phi = K * \psi$

A7: $K * (\phi \wedge \psi) \subseteq C_n((K * \phi) \cup \{\psi\})$

A8: If $\neg\psi \notin K * \phi$ then $C_n((K * \phi) \cup \{\psi\}) \subseteq K * (\phi \wedge \psi)$

The last two axioms enforce a coherence on $*$ by imposing a lower- and upper bound on the outcome. Taking $K * (\phi \wedge \psi)$ to mean the (iterated) revision $(K * \phi) * \psi$, from Axiom **A3** we then have that $K * (\phi \wedge \psi) \subseteq C_n((K * \phi) \cup \{\psi\})$, which is just the belief set formed when we first revise by ϕ and then add ψ to the outcome. This set serves as an upper bound for the outcome of the revision. Axiom **A8** is then simply a rewriting of Axiom **A4**, with the belief set $C_n((K * \phi) \cup \{\psi\})$ now serving as a lower bound for the outcome.

Axiom **A2** has been questioned in [18] ([19]) because as a consequence, an agent can no longer choose not to believe what it has just been told. A further criticism by [18] ([19]) is that, although the new information is believed by the agent, the axioms do not stipulate whether the agent then knows the new information in the sense of Section 1.1. The AGM postulates also do not restrict the structure of

ϕ beyond that it should be a (well-formed) formula of L. Thus, as in [18] ([19]) we may ask how an agent should revise its beliefs with a logical constant such as $\perp$ (falsum), or whether it should accept $\perp$ without question.

A more fundamental difficulty, questioned in works such as [18] ([19]), [25] and [11], is that a collection of beliefs must be deductively closed. This requirement raises problems of computational complexity because, for example, determining the logical consequences of a collection of beliefs may be an infinite task. It also forces us to account for all the beliefs of an agent during a revision, rather than only those relevant to the new information. The philosophical fidelity of the requirement is questioned because, for example, an agent must then know all the consequences of its beliefs – we might believe that global warming is occurring, but how could we know all the consequences of this belief?

Several contemporary approaches to belief revision seek to handle the problem of deductive closure by appeal to the notion of paraconsistency, where a logic is paraconsistent if its entailment relation $\vdash$ is not "explosive", *i.e.* for any well-formed formulae ϕ and ψ, it is not the case that $\{\phi, \neg\phi\} \vdash \psi$. An agent can then reason using smaller, locally consistent sets of beliefs. In the next section, we present such an approach to belief revision.

1.2.2 The B-Structures Approach

In this section, we describe the B-structures model of [11] in more detail, beginning with an overview of the language-splitting model of which B-structures are a natural extension. We follow directly the exposition given in [11], and most of the results, definitions and examples presented here may be found there. Our purpose here is only to present the model of [11], so definitions and results are generally introduced without further comment or proof. To avoid notational overloading in the rest of our work, we have adapted the notation used in [11] to match our conventions.

Let $L = \{p_1, p_2, \ldots, p_n, \mathsf{true}, \mathsf{false}\}$ be a finite propositional language, where the p_i are propositional atoms and true and false are logical constants. As before, let Φ_L denote the set of well-formed formulae of L formed by applying the binary connectives $\wedge$ (AND) and $\vee$ (OR) and the unary connective $\neg$ (NOT) in the standard manner. For any non-empty subset $P \subseteq L$, we shall call Φ_P a subject. Symbols such as p, q and r stand for propositional atoms, while Greek letters such as ϕ and ψ stand for arbitrary well-formed formulae in Φ_L. For a set $X \subseteq \Phi_L$ of formulae, $C_n(X)$ represents the closure of X under logical consequence. The set X is consistent if $C_n(X) \neq \Phi_L$, and a theory of L if $X = C_n(X)$. We use the letter T, possibly subscripted, to denote a theory. The revision of T with ϕ is denoted by $T * \phi$, and $T \dot{+} \phi$ is taken to mean $C_n(T \cup \{\phi\})$ where ϕ is joined to T without regard for consistency.

We begin with the idea that the language L can be split into sub-languages relative to a given theory T. Definition 1.1 may be compared to Definition 2 in [11].

Definition 1.1. Let $\mathcal{L} = \{L_1, L_2, \ldots, L_n\}$ be a family of (mutually disjoint) subsets of L, and let T be a theory of the language L. Then $L_1, \ldots, L_n$ *split* L *relative to* T if and only if for each i in $1, 2, \ldots n$ there exists $\phi_i \in \Phi_{L_i}$ such that $T = C_n(\{\phi_1, \phi_2, \ldots, \phi_n\})$. The family $\mathcal{L}$ separates T into $T_1, T_2, \ldots, T_n$ such that for each $i = 1, 2, \ldots, n$, $T_i \subseteq \Phi_{L_i}$. We call $\mathcal{L}$ a T*-splitting (of* L*)* and say that T is *generated* by the T_i expressed in each sub-language L_i. Given a language $L' \subseteq L$, we say that T *is confined to* $\Phi_{L'}$ if $T = C_n(T \cap \Phi_{L'})$.

A T-splitting $\{L_1, L_2, \ldots, L_n\}$ of L thus partitions T into $T_1, T_2, \ldots, T_n$ such that each T_i is confined to the subject Φ_{L_i} (compare also pp4–5 in [39]). It is possible for one T-splitting to refine another. The notion of refinement set out in Definition 1.2 is based on Definition 20.1 in [56].

Definition 1.2. Let $\mathcal{L}$ and $\mathcal{L}'$ be families of subsets of L such that $\bigcup \mathcal{L} = \bigcup \mathcal{L}'$. We say that $\mathcal{L}$ *refines* $\mathcal{L}'$ if for every $L \in \mathcal{L}$ there is $L' \in \mathcal{L}'$ with $L \subseteq L'$.

The extent to which one T-splitting may refine another is limited, and the existence of this limit is captured by Lemma 1.3 (compare Lemma 1 in [11], a proof of which may be found in [39]).

Lemma 1.3. *Let T be a theory of the language L. There exists a unique T-splitting $\mathcal{L}$ of L such that for every T-splitting $\mathcal{L}'$, $\mathcal{L}$ refines $\mathcal{L}'$.*

The finest T-splitting of L allows us to represent the beliefs of an agent in terms of a collection of disjoint subjects, so subjects are not combined unnecessarily. In turn, the subject matter of a belief $\phi \in \Phi_L$ is determined by the set of propositional atoms of L from which it is generated. Lemma 1.4 corresponds to Lemma 2 in [11], while Definition 1.5 is adapted from Definition 3 in [11].

Lemma 1.4. *Let ϕ be a formula of a finite propositional language L. There exists a smallest language $L_\phi \subseteq L$ with which ϕ can be expressed. That is, there is $L_\phi \subseteq L$ and a formula $\psi \in \Phi_{L_\phi}$ such that $\phi \Leftrightarrow \psi$, and for any other $L' \subseteq L$ for which there is $\psi' \in \Phi_{L'}$ such that $\phi \Leftrightarrow \psi'$, $L_\phi \subseteq L'$.*

Definition 1.5. Let ϕ be a formula of a finite propositional language L. The *subject matter of* ϕ is the smallest language, denoted L_ϕ, of a formula that can be used to express ϕ. Two formulae ϕ and ψ of L are *relevant* to each other if $L_\phi \cap L_\psi \neq \emptyset$.

Remark 1.6. The formulae $\phi = p \wedge (q \vee \neg q)$ and $\psi = r \wedge (q \vee \neg q)$ appear to be relevant to each other because they share the propositional atom q. However, they are not because $L_\phi = \{p\}$ and $L_\psi = \{r\}$. Definition 1.5 thus prevents two formulae from being relevant to each other in a trivial way. The authors of [11] concede that a sharing of propositional atoms is inadequate as a formulation of relevance because it will not capture all of the nuances of meaning that usually attend the notion.

Suppose that we have a theory T and a T-splitting $\mathcal{L} = \{L_1, L_2, \ldots, L_n\}$ of L, and that we are given the new information ϕ. Informally, to find which formulae in T are affected by ϕ, we compare L_ϕ with each L_i in $\mathcal{L}$ to determine if ϕ is relevant to T_i. The T-splitting $\mathcal{L}$, however, must be such that ϕ can be expressed by at least one sub-language in $\mathcal{L}$, which leads us to the following notion of compatibility (compare Definition 4 in [11]).

Definition 1.7. Given a theory T, a finite propositional language L and a formula ϕ of L, let L_ϕ^T be the smallest language such that $L_\phi \subseteq L_\phi^T$ and $\{L_\phi^T, L \setminus L_\phi^T\}$ is a T-splitting of L. We say that L_ϕ is *compatible* with the T-splitting $\{L_\phi^T, L \setminus L_\phi^T\}$ of L.

The language L_ϕ^T may be formed as the smallest union of elements of the finest T-splitting of L by which ϕ can be expressed. To illustrate, suppose that $T = C_n(\{p \vee q, r, s\})$ is a theory of a finite propositional language $L = \{p, q, r, s\}$, and let $\alpha = q \vee r$. Then $\{\{p, q\}, \{r\}, \{s\}\}$ is a T-splitting of L, $L_\alpha = \{q, r\}$, $L_\alpha^T = \{p, q, r\}$ and $\{\{p, q, r\}, \{s\}\}$ is a T-splitting of L with which L_α is compatible (compare the example on p5 in [11]).

The members of a T-splitting $\mathcal{L}$ of a given language L need not be disjoint since two subjects may share propositional atoms. Furthermore, two theories confined to (different) members of $\mathcal{L}$ may be consistent of their own but inconsistent when considered jointly. We then have the idea of local consistency versus global inconsistency, which brings us to the notion of a belief structure, or B-structure. A B-structure is defined in three parts, motivated by the interplay between local consistency and global inconsistency. Definitions 1.8, 1.10 and 1.11 correspond respectively to Definitions 7, 8 and 9 in [11].

Definition 1.8. A *B-structure on L* is a set $\mathcal{B} = \{(L_i, T_i)\}_{i \in I}$ where $I = \{1, 2, \ldots, n\}$, $L = \bigcup_{i \in I} L_i$ and for each $i \in I$, T_i is a consistent, finitely axiomatisable theory in L_i. For each $i \in I$, Γ_i is a set of explicit beliefs of an agent, expressed in the language L_i, such that $T_i = C_n(\Gamma_i)$.

Remark 1.9. In Definition 1.8, the Γ_i are often referred to as belief bases. A belief base is a set of beliefs that is not necessarily deductively closed, and its elements are viewed as more basic beliefs in that the full beliefs T_i of the agent about the subject Φ_{L_i} can be recovered as $C_n(\Gamma_i)$ (compare [25]). The beliefs in Γ_i and T_i are respectively called *explicit* and *implicit* beliefs.

A B-structure is a natural extension of a T-splitting of L. Each (L_i, T_i) implicitly pairs a subject covered by a particular language with a deductively closed, consistent theory about that subject. The join of the theories is not guaranteed to be consistent, however, so a B-structure captures the idea that we are "reasonably rational" agents [11]. Inconsistencies arise when for some $i, j \in I$, T_i and T_j offer divergent views on the truth of some formula $\phi \in \Phi_L$. In this case, L_i and L_j must share at least the propositional letters contained in ϕ and to cater for this overlap, [11] introduces a controlled form of overlap called a k-partition.

Definition 1.10. Let L be a finite propositional language, and let $\{L_i\}_{i \in I}$, where $I = \{1, 2, \ldots, n\}$, be a family of sub-languages of L such that $L = \bigcup_{i \in I} L_i$. Then, for $1 \leq k \leq n$, $\{L_i\}_{i \in I}$ is a *k-partition of L* if any propositional symbol $p \in L$ occurs in at most k of the sub-languages L_i.

Given the overlap determined by a k-partition, the notion of m-consistency captures the extent to which the beliefs held by an agent are coherent in the presence of global inconsistency.

Definition 1.11. Let $\mathcal{B} = \{(L_i, T_i)\}_{i \in I}$, where $I = 1, 2, \ldots, n$, be a belief structure on a finite propositional language L. Then, for $1 \leq m \leq n$, $\mathcal{B}$ is *m-consistent* if any m of the T_i are jointly consistent.

The approach to the global inconsistency taken in [11] differs from paraconsistent approaches such as in [25] because it allows full use of classical propositional logic within each subject, but resorts to a multi-valued logic when reasoning with the combined subjects. To show how an agent uses a B-structure to perform reasoning actions such as answering queries, the authors use a 4-valued logic with truth values $\mathbf{T} = \{\bot, \mathsf{true}, \mathsf{false}, \top\}$. These values indicate, respectively, a lack of information, truth, falsity and inconsistency. The assignment $v_{\mathcal{B}} : \Phi_L \longrightarrow \mathbf{T}$ described in Definition 1.12 then tells us whether a particular formula follows from a given B-structure (compare Definition 10 in [11]).

Definition 1.12. Let $\mathcal{B} = \{(L_i, T_i)\}_{i \in I}$, where $I = \{1, 2, \ldots, n\}$, be a B-structure on a finite propositional language L. Let ϕ be a formula of L and let L_ϕ be the smallest sub-language of L with which ϕ can be expressed. Let $\Gamma_\phi = \bigcup \{\Gamma_i \mid L_i \cap L_\phi \neq \emptyset\}$. We define $v_{\mathcal{B}} : \Phi_L \longrightarrow \mathbf{T}$ such that if Γ_ϕ is consistent, then

i) if $\Gamma_\phi \vdash \phi$, then $v_{\mathcal{B}}(\phi) = \mathsf{true}$,

ii) if $\Gamma_\phi \vdash \neg\phi$, then $v_{\mathcal{B}}(\phi) = \mathsf{false}$, and

iii) $v_{\mathcal{B}}(\phi) = \bot$ otherwise

and if Γ_ϕ is inconsistent, then $v_{\mathcal{B}}(\phi) = \top$.

Several possibilities for belief revision over a B-structure are given in [11]. Each method, however, retains the spirit of AGM-style belief revision, and the proposed belief revision operations are thus still modelled as functions. In particular, with what is referred to as "Option B Revision" in [11], theories that are affected by new information are merged. Thus, if ϕ is an item of new information and L_ϕ is the smallest language of ϕ, we again let $\Gamma_\phi = \bigcup\{\Gamma_i \mid L_i \cap L_\phi \neq \emptyset\}$, so that $T_\phi = C_n(\Gamma_\phi)$. We then replace all languages L_i for which $L_i \cap L_\phi \neq \emptyset$ with the combined language $\bigcup\{L_i \mid L_i \cap L_\phi \neq \emptyset\}$, and we replace the corresponding T_i with $T_\phi * \phi$, the revision of T_ϕ with ϕ. In the authors' opinion, although this approach is reasonable in most situations, it has the disadvantage of bringing about a progressive "lumping together" of the different languages, belief bases and theories (see p18 in [11]).

Consider now an agent with beliefs $\{p, q, p \wedge q\}$, where p and q are propositional atoms. If the agent is then told that $\neg p \vee \neg q$ is the case, one could reasonably expect that $p \wedge q$ is discarded. However, should the agent reject p or q? There are then two possible outcomes, depending on whether the agent discards p or q. The functional approach to belief revision does not permit such non-determinism, and in the next section, we describe a more recent approach that allows for more than one outcome by using relations to model belief revision.

1.2.3 The Relational Approach

In [36], the authors provide an AGM-style axiomatisation for relational belief revision and develop a belief revision relation using what are called epistemic entrenchments. Informally, an epistemic entrenchment captures how strongly an agent might believe a particular statement or, equivalently, how willing the agent might be to retract the statement when given new information. The exposition in [36] is fairly detailed, and we present here only the highlights of the work.

First, let us fix a sentential language $\mathcal{L}$ with atomic sentences, typically denoted by letters such as p and q, the logical constant $\perp$ (falsum), the binary logical connectives $\wedge$ and $\vee$, the unary logical connective $\neg$, possibly some non-classical n-ary logical connectives, and parentheses. The logical connectives $\rightarrow$ and $\leftrightarrow$ can be defined in terms of $\neg$, $\vee$ and $\wedge$ as with a classical propositional language, and the well-formed formulae of the language are given by

$$\phi \quad ::= \quad \perp \mid p \mid \neg\phi \mid \phi \wedge \psi \mid \phi \vee \psi$$

If c_n is an n-ary logical connective and $\phi_1, \ldots, \phi_n$ are well-formed formulae then $c_n(\phi_1, \ldots, \phi_n)$ is also a well-formed formula. The set Φ is the smallest set to contain all well-formed formulae.

A logic L of $\mathcal{L}$ is a subset of Φ that is closed under modus ponens and contains all tautologies of the language. For any set Γ of sentences and any logic L, we define a consequence operator C_n^L such that $C_n^L(\Gamma)$ is the smallest set of sentences that is closed under modus ponens and contains $\Gamma \cup L$. If $\phi \in C_n^L(\Gamma)$, we call ϕ an L-consequence of Γ and write $\Gamma \vdash_L \phi$. The logic L is consistent if $L \neq \Phi$, and Γ is L-consistent if $C_n^L(\Gamma) \neq \Phi$. A set G of sentences is an L-theory if $G = C_n^L(G)$. The set G of sentences is L-maximal if G is L-consistent and for any set H of sentences, if $G \subseteq H$ and H is L-consistent, then $G = H$. We use the letters such as G and H to denote an L-theory and $\mathcal{T}_L$ to denote the set of all such L-theories.

Next, we fix a consistent logic L of the language $\mathcal{L}$, and take the L-theory G to be the set of beliefs held by an agent. From Definition 3.1 in [36], an epistemic entrenchment for G is then a binary relation $\leq_e \subseteq \Phi \times \Phi$ such that for $\phi, \psi, \gamma \in \Phi$,

E1: If $\phi \leq_e \psi$ and $\psi \leq_e \gamma$ then $\phi \leq_e \gamma$

E2: If $\phi \vdash_L \psi$ then $\phi \leq_e \psi$

E3: If $\phi \leq_e \psi$ and $\phi \leq_e \gamma$ then $\phi \leq_e \psi \wedge \gamma$

E4: If $\perp \notin G$ then $\phi \notin G$ if and only if $\phi \leq_e \perp$

E5: If $\top \leq_e \phi$ then $\vdash_L \phi$, where $\top = \neg\bot$

Certainly, $\phi \leq_e \phi$, and together with Axiom **E1**, $\leq_e$ is then a quasi-order. The agent does not have an opinion on sentences outside of G, *i.e.* it neither believes nor disbelieves them. Under Axiom **E4**, such sentences are less entrenched than anything that the agent does believe, and in this way, $\leq_e$ may be extended to all the well-formed formulae of Φ. Correspondingly, by Axiom **E5** any theorem of the language is more entrenched than anything that the agent does believe.

Remark 1.13. Axioms **E1–E5** are adapted from the original epistemic entrenchment postulates presented in [20]. Let L be a language that is closed under application of the logical connectives $\wedge$, $\vee$, $\neg$ and $\rightarrow$, and let T be a consistent theory in L. By [20], an epistemic entrenchment $\leq_e$ relative to T is then a binary relation on the well-formed formulae of L such that

EE1: For all $\phi, \psi, \gamma \in L$, if $\phi \leq_e \psi$ and $\psi \leq_e \gamma$ then $\phi \leq_e \gamma$

EE2: For all $\phi, \psi \in L$, if $\phi \vdash_L \psi$ then $\phi \leq_e \psi$

EE3: For all $\phi, \psi \in L$, $\phi \leq_e \phi \wedge \psi$ or $\psi \leq_e \phi \wedge \psi$

EE4: If $T \neq L$ then $\phi \notin T$ if and only if $\phi \leq_e \psi$ for all $\psi \in L$

EE5: If $\phi \leq_e \psi$ for all $\phi \in L$, then $\vdash_L \psi$

An important difference between the two sets of axioms is that, from **EE3**, an L-consistent theory T is totally ordered by $\leq_e$, whereas with the **E** axioms it is not. Axiom **E2** (**EE2**) has also been questioned in [14]. There, the idea is that a conditional statement such as $\phi \vdash_L \psi$ (or equivalently $\vdash_L \phi \rightarrow \psi$) represents a "regularity" such as an "empirical law" that governs the functioning of an environment. The study reported on in [14] found that participants were more likely to disbelieve the conditional when faced with contradictory information than they were to give up belief in the antecedent or consequent. Lindström *et al.* (in [36]) and also Gärdenfors and Makinson (in [20]) seem to prefer that an agent should give up ϕ and ψ rather than $\vdash_L \phi \rightarrow \psi$.

The authors of [36] then use $\leq_e$ to define subsets of Φ called fallbacks, which effectively are formed as filters relative to $\leq_e$ (compare p44 in [13]). Fallbacks contained in G are seen as "sub-theories" of G, obtained by retracting certain formulae from G and only retaining those that are at least as entrenched as certain formulae in the fallback.

Given a formula ϕ, a fallback $H \subseteq G$ is called ϕ-permitting if $\neg\phi \notin H$. It is a maximal ϕ-permitting fallback for G if for any fallback $K \subseteq G$ with $H \subset K$, $\neg\phi \in K$. Relational revision is then defined in terms of the fallbacks of G. Informally, when the agent is given the new information ϕ, it chooses a maximal ϕ-permitting fallback of G and simply adds ϕ to it. The set of revised beliefs then contains the logical consequences of this new set. A set H of formulae is then a possible revision of G with ϕ if either $\neg\phi \in L$ and $H = \Phi$ or there exists a maximal ϕ-permitting fallback K of G such that $H = C_n^L(K \cup \{\phi\})$. The existence of such a fallback is established by appeal to Zorn's Lemma (compare the proof of Proposition 3.47).

Note that [36] does not indicate how the new information is incorporated into the epistemic entrenchment $\leq_e$. Furthermore, because we may have $\phi, \psi \in G$ such that neither $\phi \leq_e \psi$ nor $\psi \leq_e \phi$, the revision H is not expected to be unique so the operation that takes G to H given ϕ is now relational rather than functional. A belief revision relation R_ϕ that revises some theory G with the information ϕ is then taken to obey the axioms listed below (compare Definition 4.2 in [36]), for which G and H are any two L-theories and ϕ and ψ are formulae in Φ:

R1: There exists $H \in \mathcal{T}_L$ such that $H \in R_\phi(G)$

R2: If $H \in R_\phi(G)$, then $\phi \in H$

R3: If $\neg\phi \notin G$ and $H \in R_\phi(G)$, then $H = C_n^L(G \cup \{\phi\})$

R4: If $\neg\phi \notin L$ and $H \in R_\phi(G)$, then $\perp \notin H$

R5: If $\vdash_L \phi \leftrightarrow \psi$, then $H \in R_\phi(G)$ if and only if $H \in R_\psi(G)$

R6: If $H \in R_\phi(G)$ and $\neg\psi \notin H$, then $C_n^L(H \cup \{\psi\}) \in R_{\phi\wedge\psi}(G)$

R7: If $H \in R_\phi(G)$ and for all $K \in \mathcal{T}_L$ we have that if $K \in R_{\phi\vee\psi}(G)$ then $\neg\phi \notin K$, then there exists $K \in \mathcal{T}_L$ such that $K \in R_{\phi\vee\psi}(G)$ and H is given by $C_n^L(K \cup \{\phi\})$

We shall call these axioms the LR postulates. Observe that if R_ϕ determines a function, the LR postulates revert to the AGM postulates as listed in Section 1.2.1.

The authors of [36] then show that belief revision relations can be derived from an epistemic entrenchment $\leq_e$ on the set G of beliefs held by the agent. The pair $(L, \{R_\phi\}_{\phi\in\Phi})$ is called a belief revision system, and is taken to be representable if, given any L-theory G, an epistemic entrenchment $\leq_e$ on G can be recovered from it. The authors also demonstrate that not every belief revision system is representable because the axioms listed above are not strong enough to exclude belief revision relations for which the changes are not minimal (compare pp17–19 in [36]).

In this section, we presented three approaches to the problem of belief revision. We would now like to take some of these ideas forward, and in the next section we present an overview of what it is that we intend to do in this dissertation.

1.3 The Aims of Our Work

We would like to model belief revision as a relational operation that is local in scope, acting on some rather than all of the beliefs of an agent. Although the B-structures approach described in Section 1.2.2 is local in scope, belief revision over a B-structure is nonetheless modelled as a function. We therefore propose to adopt the B-structures approach of [11] and extend it to the more general case presented in [36], where the belief revision operation is modelled as a relation.

To combine these two approaches, we require a framework in which to model belief revision as a local, non-deterministic operation, and our goal in this dissertation is to develop such a framework. As we saw from Section 1.2.2, subjects play an important role in the formulation of a B-structure, so it is here that we begin our development.

1. *We take a subject S to be the smallest set to contain all of the well-formed formulae that arise from the application of a finite set of connectives to a countable set L of propositional atoms.*

For ease of exposition, the model of [11] was developed over a finite set of propositional atoms. The authors point out, however, that most of the results in [11] also hold for the countable case. We can extend their model by deriving S from countably many propositional atoms instead, since L then contains placeholders for facts about S that are known about as well as those that must still be discovered. An agent can then learn new facts about S by itself or from other agents. We can thus also extend the model of [11] by including, for any $p \in L$, the possibility that an agent does not know *about* p, so it can neither believe nor disbelieve p.

In [36], a belief revision relation was derived from a quasi-order called an epistemic entrenchment. A quasi-order on a set can be regarded as a property. Points that are "higher up" in the order then exhibit the property more strongly than those that are "lower down" (compare the positive and negative properties of Chapter 5 in [8]). By replacing an epistemic entrenchment with an unspecified quasi-order, restricted by additional conditions if necessary, we can extend the model of [36] to obtain a more general version of relational belief revision where formulae are ordered by how strongly they exhibit a given property. This leads us to the next step in our development.

2. *We use an external component to induce a quasi-order* $\leq_S \subseteq S \times S$ *on the members of* S.

This approach is more flexible than if we defined the order directly on S, because the order is then not intrinsic to S as in [36] but sits outside of it and provides a "view" or opinion of it. The structure of the external component induces the order, and we can define this structure down to arbitrary detail, independently of S. By using more than one external component, different orders can then be induced on S, and the structure of each component can vary independently from that of the others. Thus, we can again extend the model of [36] and compare formulae by how strongly they exhibit several different properties rather than just one.

A given B-structure bears a semblance of uniformity because its sub-languages all stem from the same parent language. For example, in each subject reasoning is conducted with the same logic, so each subject enjoys the same expressive and deductive capability as other subjects. It may not be appropriate for all the subjects to employ the same logic, since a particular subject might be better modelled by a different logic than the other subjects. We address this concern by developing a logic for each subject of the B-structure. In the process we are again able to extend [11]. The next step in our development is thus to derive a logic on S.

3. *We use an external component to develop a logic on* S*, and we use the logic to develop a family of belief revision relations.*

To develop the logic on S, we appeal to the idea of "preservation of degrees of truth", as described in [16]. By using the same external component as for step 2, we are able to construct a logic for which the notion of logical consequence has a pleasing interaction with the order $\leq_S$.

In a restricted sense, the B-structures model of belief revision may be considered as a top-down approach, where we seek to solve the problem of belief revision on a set T of beliefs by dividing T into smaller sets and performing revision within each set as needed. In our case, we do the reverse by taking a family of subjects, each independently developed with its own order and family of belief revision relations, and combining them into a single, unified structure. To effect the combination, we appeal to the theory of manifolds, which leads us to the next step.

4. *We use a structure called a manifold to combine the subjects into a single, unified structure.*

Manifold theory seeks to formalise how complex topological spaces can be constructed from simpler ones with well-known and accessible structure. The resulting space then "locally" resembles the simpler spaces. For manifolds, the simpler spaces are usually $\mathbb{R}^n$, which has the advantage that one can confer properties such as smoothness on a manifold.

In our case, we deal with entities such as propositions so the local structure is more likely to resemble an ordered set or a logic, and notions like smoothness may no longer be appropriate. If we are to apply manifold theory to combine the given subjects, we will need to drop the requirement that a manifold should resemble $\mathbb{R}^n$ locally. To this end, we appeal to the very general definition of a manifold in terms of sheaf theory provided in [52].

Chapter Outline:

Chapter 2: A Framework for Belief Revision. In this chapter, we develop the framework needed for our model of relational belief revision. We provide a model of a subject (step 1) and the external structure and technique by which a quasi-order is induced on it (step 2). We are then able to define the manifold that we shall use to combine a family of subjects into a single structure (step 4). The model is developed generally, without reference to any particular context, and the specialisations required for relational belief revision are provided in Chapters 3 and 4.

Chapter 3: Disposition and Logic. In Section 1.1, we described how belief can be characterised as a propositional attitude, which is a disposition towards a particular proposition or a state of affairs in which that proposition is true. Intuitively, these dispositions can be held to differing degrees, and we exploit this feature to specialise the external structures of step 2. We then develop the logic referred to in step 3.

Chapter 4: Belief Revision by Disposition. With the framework of Chapter 2 and the specialisation and logic of Chapter 3, we can now develop a family of belief revision relations in the style of [36] (step 3). To illustrate how the framework we have provided can be applied to related areas of research, we also provide a worked example that is based on an account of belief presented in [44].

In closing, we quote from [29] (page XX):

> "Abstract algebra cannot develop to its fullest extent without the infusion of topological ideas, and conversely if we do not recognise the algebraic aspects of the fundamental structures of analysis our view of them will be one-sided."

Of this quote, algebraic topology stands as an example where answers to algebraic questions are sought from a topological vantage point, while answers to topological questions are explored through algebraic images of topological spaces. It is common in mathematics to use results from one field to answer questions in another. This, together with the idea from the theory of emergence (see for example [27]) that complex and even unexpected behaviour can arise from the combination of simpler behaviours, served as primary motivation for our approach to the problem.

Work along these lines, *viz.* the deconstruction of a given structure into simpler structures and the reconstruction of a particular structure from a collection of structures (although without recourse to manifold theory), has been done in [48] to model the action of concurrent, cooperating systems. In a certain sense, the language-splitting model of [11], as it evolves under the action of some belief revision operation, can be seen as a collection of cooperating systems, and thus we are led to ask whether a similar investigation as in [48] could not also be carried out for belief revision. This brings us to the central hypothesis of our work, *viz.* the idea that

> Mathematical structures with which to study and translate between global and local behaviour can be applied to model the process of belief revision as a local, non-deterministic operation.

The development of a framework in which to model belief revision in this way is the goal and contribution of this dissertation, and provides an affirmative response in support of our hypothesis.

Chapter 2

A Framework for Belief Revision

In this chapter, we construct the framework that we will use to model belief revision as a local, non-deterministic operator. We develop the framework generally and without reference to a particular context, deferring the specialisations required for relational belief revision to Chapters 3 and 4.

The framework is presented as a manifold that is constructed from a family of topological spaces, each equipped with a quasi-order. The universe of each space is the carrier set of an algebra, and the quasi-order is induced on the space by a family of homomorphisms to a second algebra that bears a partial order. The manifold is then constructed in such a way that it inherits a quasi-order from the topological spaces in the given family.

The approach that we use to build the manifold relies on techniques and constructions from the field of sheaf theory. Our presentation of these techniques is derived from [52] and many of the definitions and results listed in this chapter may be found there, although we have adapted the notation to suit our needs better. Several concepts in sheaf theory have their origins in category theory, and although not many category theoretic concepts are used in this chapter, the interested reader may nonetheless refer to works such as [37] and [38]. We proceed as follows.

Chapter Guide:

Section 2.1: Introduction to Sheaf Theory. In this section, we present the concepts from sheaf theory that we will need to develop our framework. We introduce presheaves and sheaves as structures indexed by the open sets of a topological space. We then discuss the stalks of a presheaf, by which we can describe the behaviour of a presheaf at a point in the underlying space.

Section 2.2: Ordered Relational Algebraic Spaces. Using the work of the preceding section, we now construct a topological space that bears a quasi-order. The quasi-order is derived from an external component in the form of an algebra together with a partial order, and is conferred on the space by means of a sheaf for which the data over each open set of the topology simulates the algebraic and relational structure of the external component.

Section 2.3: Ordered Relational Algebraic Manifolds. The work of Section 2.2 allows us to construct a manifold from a family of topological spaces, each equipped with a quasi-order. The definition of a manifold that we provide is an adaptation of the sheaf-theoretic definition of a manifold given in [52], and allows the manifold to inherit an order from the family of topological spaces. We conclude by setting out a construction procedure for the manifolds we have defined.

2.1 Introduction to Sheaf Theory

Here and elsewhere, if X is a set together with a topology, we write ΩX for the topology and refer to $(X, \Omega X)$ as "the (topological) space X". We use the notation $[X \longrightarrow Y]$ to denote the family of all functions with domain X and codomain Y, and we write $f : X \longrightarrow Y : x \mapsto f(x)$ (for example, $f : \mathbb{R} \longrightarrow \mathbb{R} : x \mapsto x^2 + 3$) to indicate the domain, codomain and action of a function f. Given $f \in [X \longrightarrow Y]$, we write $f_{|U}$ to denote the restriction of f to $U \subseteq X$. Similarly, if $R \subseteq X \times X$ is a binary relation on X, we write $R_{|U}$ to mean $R \cap U \times U$, the restriction of R to U. We follow [56] and write $f(U)$ rather than $f[U]$ to mean $\{f(x) \mid x \in U\}$.

2.1.1 Presheaves and Sheaves

Definition 2.1 is based on Definition 1.1.1 and Definition 2.2 on Definitions 2.1.1, 2.1.3 and 2.1.4 in [52] (for which please compare 1.3(ii) on p171 in [29]).

Definition 2.1. Let X be a topological space. A *presheaf F of sets on X* is given by two pieces of information, *viz.*

i) For each $U \in \Omega X$, a set $F(U)$ called the *set of sections of F over U*

ii) For $U, V \in \Omega X$ with $V \subseteq U$, a restriction map $\rho^U_V : F(U) \longrightarrow F(V)$ such that

   a) The restriction map ρ^U_U is the identity map on $F(U)$

   b) For $U, V, W \in \Omega X$ with $W \subseteq V \subseteq U$, $\rho^U_W = \rho^V_W \circ \rho^U_V$

In our work, we shall use presheaves of sets on a given topological space almost exclusively , and unless otherwise indicated, any presheaf that we use may be assumed to be a presheaf of sets on the given space.

A presheaf is called a sheaf if its sections are compatible in the following sense.

Definition 2.2. Let X be a topological space, and for any $U \in \Omega X$ let $\{U_i\}_{i \in I}$ be an open cover of U with $U = \bigcup_{i \in I} U_i$. Let F be a presheaf of sets on X. We call F a *monopresheaf of sets on X* if and only if F is such that for any $s, s' \in F(U)$,

$$\forall i \in I.[\rho^U_{U_i}(s) = \rho^U_{U_i}(s')] \quad \Rightarrow \quad s = s'$$

Let $\{s_i\}_{i \in I}$ be a family of sections over the U_i with $s_i \in F(U_i)$ for each $i \in I$. If for any $i, j \in I$ we have

$$\rho^{U_i}_{U_i \cap U_j}(s_i) = \rho^{U_j}_{U_i \cap U_j}(s_j)$$

and there exists a section s in $F(U)$ such that $\rho^U_{U_i}(s) = s_i$ for each $i \in I$, then F satisfies the *glueing condition* on U. A monopresheaf of sets on X that satisfies the glueing condition on X is called a *sheaf of sets on X*.

We may also consider the restriction of a presheaf F on a space X to some $U \in \Omega X$. Definition 2.3 and Corollary 2.4 arise from Exercise 3 in Chapter 2 of [52].

Definition 2.3. Let F be a presheaf on a topological space X. For any $U \in \Omega X$, the *restriction $F_{|U}$ of F to U* is defined such that for any $V \in \Omega X$ with $V \subseteq U$, $F_{|U}(V) = F(V)$. Accordingly, the restriction maps for $F_{|U}$ are obtained by restricting the family of restriction maps for F to those maps ρ^V_W for which $V, W \in \Omega X$ and $W \subseteq V \subseteq U$.

Corollary 2.4. *If F is a presheaf (resp. sheaf) on a topological space X and $U \in \Omega X$, then $F_{|U}$ is a presheaf (resp. sheaf) on U, where U has the relative topology.*

Proof. From Definition 2.1, if F is a presheaf on X, then for any $U \in \Omega X$, $F_{|U}$ is a presheaf on U, where U has the relative topology. From Definition 2.3, if F is a sheaf on X then $F_{|U}$ satisfies the monopresheaf and glueing conditions of Definition 2.2 and hence is a sheaf on U. □

Given more than one presheaf on X, we can translate between them via presheaf morphisms. Definition 2.5 is based on Definition 1.5.1 in [52].

Definition 2.5. Let F and G be two presheaves on a topological space X. A *presheaf morphism from F to G* is a map $h : F \longrightarrow G$ that assigns to each $U \in \Omega X$ a map $h_U : F(U) \longrightarrow G(U)$ such that for any $U, V \in \Omega X$ with $V \subseteq U$,

$$h_V \circ \rho^U_V = \varrho^U_V \circ h_U$$

where ρ^U_V and ϱ^U_V are the restriction maps for F and G respectively. If $f : F \longrightarrow G$ and $g : G \longrightarrow H$ are two presheaf morphisms, their composition $g \circ f$ is defined by

$$(g \circ f)(U) = g_U \circ f_U : F(U) \longrightarrow H(U)$$

A presheaf morphism $h : F \longrightarrow G$ is *an isomorphism of presheaves* if and only if there is a presheaf morphism $g : G \longrightarrow F$ such that $h \circ g = \mathrm{id}_G$ and $g \circ h = \mathrm{id}_F$, where, if H is a presheaf on X and $U \in \Omega X$, $\mathrm{id}_H : H \longrightarrow H$ is given by $\mathrm{id}_H(U) = \mathrm{id}_{H(U)}$, the identity map on $H(U)$.

In Definition 2.5 we have used the notion that two sets A and B are isomorphic if there is a bijection between them (compare p171 in [34]). We shall use this idea in the next proposition, which is based on Proposition 1.5.2 in [52] and gives us further insight into isomorphisms of presheaves. The proof of this result is given in [52], so we will not repeat it here.

Proposition 2.6. *Let F and G be presheaves on a topological space X and let $f : F \longrightarrow G$ be a presheaf morphism from F to G. Then the following are equivalent:*

i) f is an isomorphism of presheaves

ii) for every $U \in \Omega X$, $f_U : F(U) \longrightarrow G(U)$ is bijective

iii) for every $U \in \Omega X$, $f_U : F(U) \longrightarrow G(U)$ is an isomorphism

In the more general case, we translate between presheaves on two different spaces. As before, the morphism should respect the restriction maps. One difference is the direction of the morphism – for two presheaves on different spaces the morphism operates from the codomain to the domain of the continuous function. The reason for this is that an open set in the domain of the function does not necessarily map to an open set in its codomain.

We first consider the image of a presheaf under the action of a continuous function. Definition 2.7 is adapted from the construction numbered 3.7.1 in [52].

Definition 2.7. Let X and Y be topological spaces, let F be a presheaf on X and let $k : X \longrightarrow Y$ be a continuous function. The *direct image of F by k* is the presheaf k_*F on Y obtained by setting

$$\begin{aligned} (k_*F)(U) &= F(k^{\leftarrow}(U)) && (U \in \Omega Y) \\ \varrho^U_V &= \rho^{k^{\leftarrow}(U)}_{k^{\leftarrow}(V)} && (U, V \in \Omega Y, V \subseteq U) \end{aligned}$$

where ρ and ϱ are the restriction maps for F and k_*F respectively.

As may be expected, if $f : F \longrightarrow G$ is a morphism of presheaves on X, then a corresponding morphism $k_*f : k_*F \longrightarrow k_*G$ is induced on Y by k. Proposition 2.8 tells us that the structure on the sets of sections of a sheaf is not affected by the formation of a direct image. It is reproduced from Proposition 3.7.3 in [52], and is readily established. We have set the proof out in more detail than was provided in [52].

Proposition 2.8. *Let X and Y be topological spaces, and let $k : X \longrightarrow Y$ be a continuous function. If F is a sheaf on X, then k_*F is a sheaf on Y.*

Proof. Let F be a sheaf on X. For $V \in \Omega Y$, let $\{V_i\}_{i\in I}$ be an open cover of V with $V = \bigcup_{i\in I} V_i$. Set $U = k^{\leftarrow}(V)$ and $U_i = k^{\leftarrow}(V_i)$, so that $\{U_i\}_{i\in I}$ is an open cover of U with $U = \bigcup_{i\in I} U_i$. Let $s, t \in F(U)$. Then because F is a monopresheaf,

$$\forall i \in I.[\rho^U_{U_i}(s) = \rho^U_{U_i}(t)] \quad \Rightarrow \quad s = t$$

where ρ is the restriction map for F. From Definition 2.7, for any $U, V \in \Omega Y$ with $V \subseteq U$ we have $\varrho^U_V = \rho^{k^{\leftarrow}(U)}_{k^{\leftarrow}(V)}$ where ϱ is the restriction map for k_*F. Hence, noting that $s, t \in k_*F(V)$,

$$\forall i \in I.[\rho^U_{U_i}(s) = \rho^U_{U_i}(t)] \quad \Leftrightarrow \quad \forall i \in I.[\varrho^V_{V_i}(s) = \varrho^V_{V_i}(t)]$$

from which we have

$$\forall i \in I.[\varrho^V_{V_i}(s) = \varrho^V_{V_i}(t)] \quad \Rightarrow \quad s = t$$

so that k_*F is a monopresheaf (Definition 2.2). Satisfaction of the glueing condition of Definition 2.2 by k_*F follows similarly, and hence k_*F is a sheaf on Y. □

A presheaf morphism that takes place relative to a continuous function k is given in terms of a k-morphism. Definition 2.9 is adapted from Definition 3.7.8 in [52], which was specialised to presheaves of abelian groups.

Definition 2.9. Let X and Y be topological spaces, let $k : X \longrightarrow Y$ be a continuous function, and let F and G be presheaves on X and Y respectively. A *morphism* $f : G \longrightarrow F$ *relative to* k (or simply, a k*-morphism*) is given by a collection of maps $f(U, V) : G(V) \longrightarrow F(U)$ for any $U \in \Omega X$ and $V \in \Omega Y$ with $U \subseteq k^{\leftarrow}(V)$ and subject to the condition that for any $U' \in \Omega X$ and $V' \in \Omega Y$ with $U' \subseteq U$ and $V' \subseteq V$ we have

$$\rho^U_{U'} \circ f(U, V) = f(U', V') \circ \varrho^V_{V'}$$

where ρ and ϱ denote the restriction maps for F and G respectively.

Intuitively, a k-morphism $f : G \longrightarrow F$ is determined by $f(k^{\leftarrow}(V), V)$ for any $V \in \Omega Y$ (compare Proposition 3.7.10 in [52]), because by setting $V' = V$ in Definition 2.9, for any $U \subseteq k^{\leftarrow}(V)$ we get

$$\rho^{k^{\leftarrow}(V)}_U \circ f(k^{\leftarrow}(V), V) = f(U, V)$$

From Definition 2.5, the family of maps of the form $f(k^{\leftarrow}(V), V) : G(V) \longrightarrow F(k^{\leftarrow}(V))$ determines a presheaf morphism $f_0 : G \longrightarrow k_*F$, so we are justified in thinking of a k-morphism from G to F as a presheaf morphism from G to k_*F. By Proposition 3.7.10 in [52], this presheaf morphism is uniquely determined by the given k-morphism. As may be expected, a k-morphism $f : G \longrightarrow F$ is an isomorphism of presheaves if k is a homeomorphism and $f_0 : G \longrightarrow k_*F$ is an isomorphism of presheaves (compare Proposition 4.1.7 in [52]).

Remark 2.10. Given presheaves F and G on the topological spaces X and Y respectively, a k-morphism from G to F is sometimes expressed as a pair (k, h) in which $k : X \longrightarrow Y$ is a continuous function and h is a presheaf morphism from G to k_*F (compare p132 in [40]). For further information, please also compare paragraph 1.8 on p174 in [29], Section 3.7 and Proposition 4.1.7 in [52] and pp1425-1427 in [28].

We may also consider the local structure of a presheaf, in the following sense. From 9K in [56] and p24 in [34], we have the following.

Definition 2.11. Let X and Y be topological spaces. A continuous map $k : X \longrightarrow Y$ is a *local homeomorphism* if every $x \in X$ has a neighbourhood U such that $k(U) \in \Omega Y$ and $k_{|U} : U \longrightarrow k(U)$ is a homeomorphism.

Definition 2.12 may be compared to Definition 4.3.6 in [52].

Definition 2.12. Let F and G be presheaves on the topological spaces X and Y respectively, and let $k : X \longrightarrow Y$ be a local homeomorphism. Then F and G are *locally isomorphic as presheaves* if every $x \in X$ has an open neighbourhood U such that

i) $k(U) \in \Omega Y$ and $k_{|U} : U \longrightarrow k(U)$ is a homeomorphism

ii) There exists an isomorphism $f_0 : G_{|k(U)} \longrightarrow (k_{|U})_*(F_{|U})$ of presheaves.

A k-morphism $f : G \longrightarrow F$ is a *local isomorphism of presheaves* if $k : X \longrightarrow Y$ is a local homeomorphism and whenever $U \in \Omega X$ is such that $k(U) \in \Omega Y$ and $k_{|U} : U \longrightarrow k(U)$ is a homeomorphism, $(f_U)_* : G(U) \longrightarrow k_*F(U)$ is an isomorphism.

In this section, we introduced presheaves and sheaves as structures that are indexed by the open sets of a topological space. To examine the data provided by the presheaf at a point in the space, we need some means of describing the behaviour of the presheaf at that point. To this end, in the next section we introduce the notion of the stalk of a presheaf.

2.1.2 The Stalks of a Presheaf

Informally, given a presheaf F on a topological space X, the behaviour of F at a point $x \in X$ can be isolated by examining sufficiently small neighbourhoods of x. The idea is that if the neighbourhood of x is small enough, the behaviour of F on the neighbourhood should be the same as its behaviour at the point. This leads us to the idea that we should determine some form of "limit" of F as the neighbourhoods around x become smaller and smaller.

To set this limit up, we begin with the notion of a directed set (compare Section 1.3 in [52]). Given an ordered set $(P, \leq)$, we denote the sets of lower and upper bounds in P of Q respectively by Q^l and Q^u, for any $Q \subseteq P$. From [13] we have the following (Definition 7.7; compare Definition 11.1 in [56]).

Definition 2.13. Let Q be a non-empty subset of an ordered set P. Then Q is said to be *directed* if, for every $x, y \in Q$ there is $z \in Q$ with $z \in \{x, y\}^u$. Equivalently, Q is directed if and only if for every finite subset F of Q there is $z \in Q$ such that $z \in F^u$.

From Definition 1.3.1 in [52], we define a directed system of sets as follows.

Definition 2.14. A *direct system of sets* is a pair $\mathbf{U} = (\mathcal{U}, \mathcal{F})$ where $\mathcal{U} = \{U_i\}_{i \in I}$ is a family of sets indexed by a directed set $(I, \leq)$ and $\mathcal{F}$ is a set of maps of the form $f_{ij} : U_i \longrightarrow U_j$, with $f_{ij} \in \mathcal{F}$ if and only if $i \leq j$ in I. The maps f_{ij} in $\mathcal{F}$ satisfy the conditions

i) for any $i \in I$, $f_{ii} = \mathrm{id}_{U_i}$, the identity function on U_i

ii) for any $i, j, k \in I$ with $i \leq j \leq k$, $f_{ik} = f_{jk} \circ f_{ij}$

Definition 2.15 is based on Definition 1.3.4 in [52].

Definition 2.15. Let $\mathbf{U} = (\mathcal{U}, \mathcal{F})$ be a direct system of sets, with $\mathcal{U} = \{U_i\}_{i \in I}$. A *target for* $\mathbf{U}$ is a pair $(V, \{g_i : U_i \longrightarrow V\}_{i \in I})$ where V is a set and for any $i, j \in I$ with $i \leq j$, $g_i = g_j \circ f_{ij}$. A *direct limit for* $\mathbf{U}$ is a target $(X, \{h_i : U_i \longrightarrow X\}_{i \in I})$ for $\mathbf{U}$ satisfying the universal property that for any target $(V, \{g_i : U_i \longrightarrow V\}_{i \in I})$ for $\mathbf{U}$ there is a unique map $f_{XV} : X \longrightarrow V$ such that for any $i \in I$, $g_i = f_{XV} \circ h_i$.

A direct limit may be characterised as follows (Theorem 1.3.8 in [52]).

Proposition 2.16. *Let $\mathbf{U} = (\mathcal{U}, \mathcal{F})$ be a direct system of sets, with $\mathcal{U} = \{U_i\}_{i \in I}$ and $\mathcal{F} = \{f_{ij} : U_i \longrightarrow U_j\}$ where $f_{ij} \in \mathcal{F}$ if and only if $i \leq j$ in I. Let $\mathbf{V} = (V, \{g_i : U_i \longrightarrow V\}_{i \in I})$ be a target for $\mathbf{U}$ such that*

i) for any $v \in V$ there is $i \in I$ such that $v \in g_i(U_i)$

ii) if $i, j \in I$ and $u_i \in U_i$ and $u_j \in U_j$, then

$$g_i(u_i) = g_j(u_j) \text{ if and only if } \exists k \in I.[i, j \leq k \text{ and } f_{ik}(u_i) = f_{jk}(u_j)]$$

Then $\mathbf{V}$ is a direct limit for $\mathbf{U}$.

It can be shown that any two direct limits of a direct system are isomorphic, and this is done in Proposition 1.3.6 in [52]. We can therefore speak of *the* direct limit of a direct system, and we denote this limit by $\varinjlim_{i \in I} U_i$. It can also be shown that every direct system of sets has a direct limit (compare Theorem 1.3.10 in [52]).

Given a presheaf F on a set X, fix $x \in X$ and let $\mathcal{U}_x = \{U \in \Omega X \mid x \in U\}$. The pair

$$\mathbf{U} = (\{F(U)\}_{U \in \mathcal{U}_x}, \{\rho^U_V : F(U) \longrightarrow F(V) \mid U, V \in \mathcal{U}_x \text{ and } V \subseteq U\})$$

is then a direct system of sets, since for each $U \in \mathcal{U}_x$, $\rho^U_U = \mathrm{id}_U$ and for $U, V, W \in \mathcal{U}_x$ with $W \subseteq V \subseteq U$, $\rho^U_W = \rho^V_W \circ \rho^U_V$. From this observation, we may make the following definition, based on Definition 1.4.1 [52] and paragraph 1.5 on p172 in [29].

Definition 2.17. Let X be a topological space, and let F be a presheaf on X. The *stalk of* F at a point $x \in X$ is the direct limit

$$F_x = \varinjlim_{U \ni x} F(U)$$

of the sets $F(U)$ as U ranges over the open neighbourhoods of x. If $s \in F(U)$ for some neighbourhood U of x, we write s_x for the image of s in F_x and call it the *germ of* s *at* x. For any $U \in \Omega X$ and $s \in F(U)$, we write $[s]$ to mean the set $\{s_x \mid x \in U\}$.

Propositions 2.18 and 2.19 correspond to Propositions 1.4.2 and 2.3.1 in [52]. The proofs are supplied in [52], so we will not reproduce them here.

Proposition 2.18. *Let F be a presheaf on a topological space X. Then,*

i) Each germ $t \in F_x$ arises as $t = s_x$ for some $s \in F(U)$, where U is an open neighbourhood of x

ii) For two germs $s_x, t_x \in F_x$, (with $s \in F(U)$ and $t \in F(V)$ say),

$$s_x = t_x \quad \Leftrightarrow \quad \exists W \in \Omega X \mid W \subseteq U \cap V.[\rho^U_W(s) = \rho^V_W(t)]$$

Proposition 2.19. *Let F be a sheaf on a topological space X. Then for any $U \in \Omega X$ and $s, t \in F(U)$,*

$$s = t \quad \Leftrightarrow \quad \forall x \in U.[s_x = t_x]$$

2.2 Ordered Relational Algebraic Spaces

In this section, we develop a topological space that is equipped with a quasi-order. Sets that carry an algebraic and relational structure play a central role in the derivation of this order, and it is with these that we begin our development.

2.2.1 Relational Algebraic Structures

We shall use the following notion of an algebra. Definitions 2.20 and 2.21 are based on Definitions 1.2, 1.3, 2.1 and 6.1 in Chapter II of [10].

Definition 2.20. A *language* or *type* of algebras is a non-empty set $\mathcal{F}$ of function symbols. Each member f of $\mathcal{F}$ is assigned a non-negative integer n called the *arity* or *rank* of f, and f is called an *n-ary function symbol*. The set of all n-ary function symbols in $\mathcal{F}$ is denoted by $\mathcal{F}_n$. A function symbol with arity 0 is called a *nullary function symbol*.

Definition 2.21. Let $\mathcal{F}$ be a language of algebras.

1. An *algebra of type* $\mathcal{F}$ is a pair $\mathbf{A} = (A, F)$ where A is a non-empty set and F is a family of operations on A indexed by the language $\mathcal{F}$ such that for each n-ary function symbol f in $\mathcal{F}$ there is an n-ary operation $f^{\mathbf{A}} : A^n \longrightarrow A$ on A. The set A is called the *universe*, *domain* or *carrier set* of $\mathbf{A}$, and each operation $f^{\mathbf{A}}$ in F is called a *fundamental operation* of $\mathbf{A}$.

2. If $\mathbf{A}$ and $\mathbf{B}$ are two algebras of type $\mathcal{F}$, a function $g : A \longrightarrow B$ is a *homomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ if for each n-ary function symbol $f \in \mathcal{F}$ and $x_1, x_2, \ldots, x_n \in A$ we have

$$g(f^{\mathbf{A}}(x_1, x_2, \ldots, x_n)) \quad = \quad f^{\mathbf{B}}(g(x_1), g(x_2), \ldots, g(x_n))$$

   If g is surjective, then $\mathbf{B}$ is called a *homomorphic image of* $\mathbf{A}$, and g is called an *epimorphism*. If g is bijective (injective), then g is an *isomorphism (embedding) from* $\mathbf{A}$ *to* $\mathbf{B}$.

A nullary function symbol gives rise to a nullary operation. If $f^{\mathbf{A}} : A^n \longrightarrow A$ is a nullary operation ($n = 0$) it is customary to define $A^0 = \{\emptyset\}$, in which case the operation is written as $f^{\mathbf{A}}(\emptyset)$ and is thought of as an element of A (compare Definition 1.1 in Chapter II of [10]). We shall adopt the convention that nullary functions are preserved by homomorphisms. That is, if $\mathbf{A}$ and $\mathbf{B}$ are two algebras of type $\mathcal{F}$, $f \in \mathcal{F}$ is a nullary operation and $g : A \longrightarrow B$ is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$, then $g(f^{\mathbf{A}}(\emptyset)) = f^{\mathbf{B}}(\emptyset)$.

To simplify our notation, we shall write f for $f^{\mathbf{A}}$ when doing so causes no confusion. Some texts prefer to write $\mathcal{F}$ as a list of arities so that, for example, $\mathcal{F} = \{2, 2, 1\}$ signifies an algebra with three function symbols of arity 2, 2 and 1 (see Example 3.3(iii) on p68 in [12] for an example of where this notation is used). We will sometimes use notation such as $\mathcal{F} = \{+, *, -\} = \{2, 2, 1\}$ to indicate that $+$ and $*$ are binary function symbols, while $-$ is a unary function symbol.

We shall use the following notion of a relational structure. The definitions of isomorphism and homomorphism given here are adaptations of Definitions 2.1 and 6.1 of Chapter II in [10] to the relational case.

Definition 2.22. Let A be a non-empty set.

1. A *relational structure* is a pair $\mathbf{A} = (A, R_A)$ where R_A is a binary relation over A.

2. If $\mathbf{A} = (A, R_A)$ and $\mathbf{B} = (B, R_B)$ are relational structures, a function $g : A \longrightarrow B$ is a *homomorphism from* $\mathbf{A}$ *to* $\mathbf{B}$ if for any $a, b \in A$,

$$aR_Ab \quad \Rightarrow \quad g(a)R_Bg(b)$$

If g is surjective, then $\mathbf{B}$ is called a *homomorphic image of* $\mathbf{A}$, and g is called an *epimorphism*. If g is bijective (injective), then g is an *isomorphism (embedding) from* $\mathbf{A}$ *to* $\mathbf{B}$ if for any $a, b \in A$,

$$aR_Ab \quad \Leftrightarrow \quad g(a)R_Bg(b)$$

We shall also be interested in sets that carry both an algebraic and relational structure.

Definition 2.23. Let $\mathcal{F}$ be a language of algebras.

1. A *relational algebraic structure of type* $\mathcal{F}$ is a triple $\mathbf{A} = (A, F, R_A)$, where (A, F) is an algebra of type $\mathcal{F}$ and (A, R_A) is a relational structure.

2. If $\mathbf{A} = (A, F, R_A)$ and $\mathbf{B} = (B, G, R_B)$ are relational algebraic structures of type $\mathcal{F}$, then a function $g : A \longrightarrow B$ is a *homomorphism (epimorphism, isomorphism, embedding) from* $\mathbf{A}$ *to* $\mathbf{B}$ if it is a homomorphism (epimorphism, isomorphism, embedding) from (A, F) to (B, G) and from (A, R_A) to (B, R_B).

To simplify our exposition, we shall abbreviate the term "relational algebraic structure of type $\mathcal{F}$" to "r-algebra (of type $\mathcal{F}$)", taking the language to be $\mathcal{F}$ if this causes no confusion. Thus, where more than one r-algebra is listed, they may all be assumed to be of the same type $\mathcal{F}$ if their types are not explicitly mentioned.

Following Definition 1.1.1 in [52], a topological space can be equipped with an r-algebraic structure by means of a sheaf for which

i) for each open set, the corresponding set of sections has an r-algebraic structure, and

ii) the restriction maps are homomorphisms of r-algebraic structures

To equip the given space with the structure of a particular r-algebra, we specialise this idea by appealing to the notion of an r-algebra simulation. Definition 2.24 may be compared to Definition 4.1.1 in [52].

Definition 2.24. Let $\mathbf{A}$ be an r-algebra.

1. An $\mathbf{A}$*-simulation* is a pair $(\mathbf{B}, f)$ where $\mathbf{B}$ is an r-algebra and $f : A \longrightarrow B$ is a homomorphism from $\mathbf{A}$ to $\mathbf{B}$ called the *structure map*.

2. If $(\mathbf{B}, f)$ and $(\mathbf{C}, g)$ are $\mathbf{A}$-simulations, a function $h : \mathbf{B} \longrightarrow \mathbf{C}$ is a *morphism of* $\mathbf{A}$*-simulations from* $\mathbf{B}$ *to* $\mathbf{C}$ if h is a homomorphism from $\mathbf{B}$ to $\mathbf{C}$ such that $g = h \circ f$.

To illustrate, let $\mathbf{A} = (A, \leq_A)$ be a partially ordered set that is also a lattice. We may then write $\mathbf{A}$ as the r-algebra $(A, \wedge^{\mathbf{A}}, \vee^{\mathbf{A}}, \leq_A)$, of type $\{\wedge, \vee\}$. Let U be a non-empty set, and consider the family $F = [U \longrightarrow A]$ of functions from U to A. Setting $\mathbf{F} = (F, \wedge^{\mathbf{F}}, \vee^{\mathbf{F}}, \leq_F)$, we define $\wedge^{\mathbf{F}}, \vee^{\mathbf{F}}$ and $\leq_F$ for $\mathbf{F}$ in terms of $\wedge^{\mathbf{A}}, \vee^{\mathbf{A}}$ and $\leq_A$ such that for any $r, s, t \in F$,

$$t = r \wedge^{\mathbf{F}} s \quad \text{if and only if} \quad \forall x \in U.[t(x) = r(x) \wedge^{\mathbf{A}} s(x)]$$

$$t = r \vee^{\mathbf{F}} s \quad \text{if and only if} \quad \forall x \in U.[t(x) = r(x) \vee^{\mathbf{A}} s(x)]$$
$$r \leq_F s \quad \text{if and only if} \quad \forall x \in U.[r(x) \leq_A s(x)]$$

Then $\mathbf{F} = (F, \wedge^{\mathbf{F}}, \vee^{\mathbf{F}}, \leq_F)$ is an r-algebra of type $\{\wedge, \vee\}$, inheriting this property via the pointwise application of $\wedge^{\mathbf{A}}, \vee^{\mathbf{A}}$ and $\leq_A$. The map $f : A \longrightarrow F : a \mapsto h$, where $h(x) = a$ for all $x \in U$, is a homomorphism from $\mathbf{A}$ into $\mathbf{F}$, and $(\mathbf{F}, f)$ is then an $\mathbf{A}$-simulation. We shall develop this example further in Section 2.2.2 (see Proposition 2.29).

Intuitively, if $\mathbf{A}$ and $\mathbf{B}$ are r-algebras and $(\mathbf{C}, f)$ is an $\mathbf{A}$-simulation, then if there is a homomorphism $h : A \longrightarrow B$ from $\mathbf{A}$ to $\mathbf{B}$, $(\mathbf{C}, f)$ can be expressed as a $\mathbf{B}$-simulation $(\mathbf{C}, g)$. All we require is an epimorphism $h_0 : B \longrightarrow h(A)$ from $\mathbf{B}$ to $\mathbf{B}_{|h(A)}$ such that $h_{0|h(A)} = \mathrm{id}_{h(A)}$, together with a homomorphism $g_0 : h(A) \longrightarrow C$ from $\mathbf{B}_{|h(A)}$ to $\mathbf{C}$ such that $f = g_0 \circ h$. From Theorem 6.5 in [10], the composition of two homomorphisms is again a homomorphism, so the function $g : B \longrightarrow C$ with $g = g_0 \circ h_0$ is a homomorphism from $\mathbf{B}$ to $\mathbf{C}$ and $(\mathbf{C}, g)$ is then a $\mathbf{B}$-simulation. We will exploit this idea in Section 2.3.2.

From Definitions 1.1.1, 2.1.4 and 1.5.1 in [52], we then have the following.

Definition 2.25. Let $\mathbf{A}$ be an r-algebra and X a topological space.

1. A *presheaf of* $\mathbf{A}$*-simulations over* X is a presheaf F of sets on X such that
   i) For each $U \in \Omega X$, $F(U)$ is an $\mathbf{A}$-simulation
   ii) For each $U, V \in \Omega X$ with $V \subseteq U$, the restriction map $\rho^U_V : F(U) \longrightarrow F(V)$ is a morphism of $\mathbf{A}$-simulations from $F(U)$ to $F(V)$

   A *sheaf of* $\mathbf{A}$*-simulations on* X is a presheaf of $\mathbf{A}$-simulations over X that satisfies the mono-presheaf and glueing conditions of Definition 2.2.

2. If F and G are two presheaves of $\mathbf{A}$-simulations over X then a *morphism of presheaves of* $\mathbf{A}$*-simulations from* F *to* G is a presheaf morphism $f : F \longrightarrow G$ such that for each $U \in \Omega X$, $f_U : F(U) \longrightarrow G(U)$ is a morphism of $\mathbf{A}$-simulations from $F(U)$ to $G(U)$.

Intuitively, if F is a presheaf (sheaf) of $\mathbf{A}$-simulations over X, then for any $U \in \Omega X$, $F_{|U}$ is a presheaf (sheaf) of $\mathbf{A}$-simulations over U, where U has the relative topology.

Given an r-algebra $\mathbf{A}$, a direct system of $\mathbf{A}$-simulations is a direct system

$$\mathbf{U} = (\{\mathbf{U}_i\}_{i \in I}, \{f_{ij} : U_i \longrightarrow U_j\}_{i,j \in I})$$

in which each $\mathbf{U}_i$ is an $\mathbf{A}$-simulation and all of the f_{ij} are morphisms of $\mathbf{A}$-simulations (compare Definition 2.14, and also Definition 1.3.11 in [52]). A target for $\mathbf{U}$ is a pair $(\mathbf{V}, \{g_i : U_i \longrightarrow V\}_{i \in I})$ such that $\mathbf{V}$ is an $\mathbf{A}$-simulation and each g_i is a morphism of $\mathbf{A}$-simulations. The notion of a direct limit for a direct system of $\mathbf{A}$-simulations may be phrased similarly (compare Definition 2.15, and also Definition 1.3.14 and Remark 1.3.15 in [52]).

As before, two direct limits for a direct system of $\mathbf{A}$-simulations are naturally isomorphic (compare Proposition 1.3.16 in [52]), and may be characterised similarly to Proposition 2.16 (compare Theorem 1.3.18 in [52]). Furthermore, it can be shown that any direct system of $\mathbf{A}$-simulations has a direct limit (compare the construction 1.3.19 and Theorem 1.3.20 in [52]), from which we deduce that the stalk of a presheaf of $\mathbf{A}$-simulations at a given point is also an $\mathbf{A}$-simulation. This then establishes the following result.

Proposition 2.26. *Let* $\mathbf{A}$ *be an r-algebra, let* X *be a topological space and let* F *be a presheaf of* $\mathbf{A}$*-simulations over* X*. For any* $x \in X$*, the stalk* F_x *of* F *at* x *is an* $\mathbf{A}$*-simulation.*

For Definition 2.27, we have adapted Definitions 4.1.2 and 4.1.6 and Proposition 4.1.7 in [52] from rings to the case of $\mathbf{A}$-simulations. In Definition 2.27(2), we have used the idea that a k-morphism $g : G \longrightarrow F$ relative to a continuous function $k : X \longrightarrow Y$, where F and G are presheaves on the topological spaces X and Y respectively, may be thought of as a morphism $g_0 : G \longrightarrow k_*F$ of presheaves on Y (compare the discussion following Definition 2.9).

Definition 2.27. Let $\mathbf{A}$ be an r-algebra of type $\mathcal{F}$.

1. A *relational algebraic space of type $\mathcal{F}$ over $\mathbf{A}$* is a pair (X, H) where X is a topological space and H is a sheaf of $\mathbf{A}$-simulations over X. The sheaf H is called the *structure sheaf* of (X, H).

2. A morphism $f : (X, G) \longrightarrow (Y, H)$ of relational algebraic spaces of type $\mathcal{F}$ over $\mathbf{A}$ is given by a continuous function $k : X \longrightarrow Y$ together with a k-morphism $g : H \longrightarrow G$ of sheaves of $\mathbf{A}$-simulations, or (equivalently) a morphism $g_0 : H \longrightarrow k_*G$ of sheaves of $\mathbf{A}$-simulations on Y. The morphism f is an *isomorphism of relational algebraic spaces of type $\mathcal{F}$ over $\mathbf{A}$*, and (X, G) and (Y, H) are *isomorphic as relational algebraic spaces of type $\mathcal{F}$ over $\mathbf{A}$* if k is a homeomorphism and g (equivalently g_0) is an isomorphism of sheaves of $\mathbf{A}$-simulations.

As before, we shall abbreviate "relational algebraic space of type $\mathcal{F}$" to "r-algebraic space (of type $\mathcal{F}$)", omitting the type $\mathcal{F}$ of the underlying r-algebra $\mathbf{A}$ if this causes no confusion. If more than one r-algebraic space is listed without mention of their respective types, all of the spaces may be assumed to be of type $\mathcal{F}$. By abuse of terminology, we shall usually write "X is an r-algebraic space (over $\mathbf{A}$)".

In the next section, we show how, by appeal to sheaf theory, we can construct an order and a topology on a set from an r-algebra.

2.2.2 Ordered Spaces over a Relational Algebraic Structure

In this section, we introduce the idea of an ordered r-algebra of type $\mathcal{F}$, which is an r-algebra in which the relation on the carrier set is a quasi-order. This order is then used to induce a quasi-order on a given topological space via the $\mathbf{A}$-simulations of Definition 2.24. We then extend the derivation of the order to the case where we have a set with only an algebraic structure, and we are then able to derive both an order and a topology on the set.

The topological structure is derived so as to have a well-defined interaction with the order. This interaction is necessary for the work we shall need to complete in Section 2.3. In turn, the order is needed for us to equip the given set with a logic, which we shall do in Section 3.2.

We begin with the idea of an ordered r-algebra.

Definition 2.28. Let $\mathcal{F}$ be a language of algebras.

1. An *ordered r-algebra of type $\mathcal{F}$* is an r-algebra $(A, F, \leq_A)$ of type $\mathcal{F}$, where $\leq_A$ is a quasi-order on A. A *bounded r-algebra of type $\mathcal{F}$* is a structure $(A, F, \mathbf{0}_A, \mathbf{1}_A, \leq_A)$ where $(A, F, \leq_A)$ is an ordered r-algebra of type $\mathcal{F}$ and $\mathbf{0}_A, \mathbf{1}_A \in A$ are two constants such that $\mathbf{0}_A \leq_A x \leq \mathbf{1}_A$ for any $x \in A$. The element $\mathbf{1}_A$ is called *top*, and $\mathbf{0}_A$ is called *bottom*.

2. If $\mathbf{A} = (A, F, \mathbf{0}_A, \mathbf{1}_A, \leq_A)$ and $\mathbf{B} = (B, F, \mathbf{0}_B, \mathbf{1}_B, \leq_B)$ are bounded r-algebras of type $\mathcal{F}$, a function $f : A \longrightarrow B$ is a *homomorphism (epimorphism, isomorphism, embedding) from $\mathbf{A}$ to $\mathbf{B}$* if it is a homomorphism (epimorphism, isomorphism, embedding) from $(A, F, \leq_A)$ to $(B, F, \leq_B)$ such that $f(\mathbf{0}_A) = \mathbf{0}_B$ and $f(\mathbf{1}_A) = \mathbf{1}_B$.

Definitions 2.24, 2.25 and Proposition 2.26 may be extended to the case of a bounded r-algebra. All that is required is that the applicable morphisms should additionally preserve the top and bottom elements. In view of Definition 2.20, we shall treat the constants $\mathbf{0}_A$ and $\mathbf{1}_A$ in a bounded r-algebra $\mathbf{A} = (A, F, \mathbf{0}_A, \mathbf{1}_A, \leq_A)$ as the constants $\mathbf{0}^{\mathbf{A}}$ and $\mathbf{1}^{\mathbf{A}}$ that arise from nullary function symbols $\mathbf{0}, \mathbf{1} \in \mathcal{F}$. By the convention established after Definition 2.21, preservation of the top and bottom elements of a bounded r-algebra by the morphisms of Definitions 2.24 and 2.25 is then standard.

Let us now presuppose a bounded r-algebra $\mathbf{A} = (A, F, \mathbf{0}_A, \mathbf{1}_A, \leq_A)$ of type $\mathcal{F}$ in which $\leq_A$ is a partial order.

Proposition 2.29. *Let U be a non-empty set, and let $\mathbf{U} = ([U \longrightarrow A], F, \leq_U)$ be an r-algebra of type $\mathcal{F}$ such that*

i) For any nullary function symbol $f \in \mathcal{F}$, if $f^{\mathbf{A}} = a_f$ then $f^{\mathbf{U}} = v_f$, where for all $x \in U$, $v_f(x) = a_f$.

ii) For any n-ary function symbol $f \in \mathcal{F}$, for $n \geq 1$,

$$v = f^{\mathbf{U}}(v_1, v_2 \ldots, v_n) \;\; \Leftrightarrow \forall x \in U.[v(x) = f^{\mathbf{A}}(v_1(x), v_2(x) \ldots, v_n(x))]$$

where $v_1, v_2 \ldots, v_n \in [U \longrightarrow A]$.

iii) For all $u, v \in [U \longrightarrow A]$,

$$u \leq_U v \quad \Leftrightarrow \quad \forall x \in U.[u(x) \leq_A v(x)]$$

and let $f : \mathbf{A} \longrightarrow \mathbf{U} : a \mapsto v$, where $v(x) = a$ for all $x \in U$. Then $([U \longrightarrow A], F, v_{\mathbf{0}}, v_{\mathbf{1}}, \leq_U)$ is a bounded r-algebra of type $\mathcal{F}$, and $(\mathbf{U}, f)$ is an $\mathbf{A}$-simulation.

Proof. From Definition 2.28, $\mathbf{U}$ is an ordered r-algebra of type $\mathcal{F}$. For $\leq_U$ as defined, it follows that $v_{\mathbf{0}} \leq_U u \leq_U v_{\mathbf{1}}$ for any $u \in [U \longrightarrow A]$. Hence $\mathbf{U}$ is bounded and so $([U \longrightarrow A], F, v_{\mathbf{0}}, v_{\mathbf{1}}, \leq_U)$ is a bounded r-algebra of type $\mathcal{F}$. Finally, the given function f defines a homomorphism of r-algebras from $\mathbf{A}$ to $\mathbf{U}$, so that $(\mathbf{U}, f)$ is an $\mathbf{A}$-simulation. □

Now let (X, H) be an r-algebraic space over $\mathbf{A}$ such that for each $U \in \Omega X$, $H(U) = [U \longrightarrow A]$. Take any $x \in X$ and $U \in \Omega X$ with $x \in U$. Then $H(U)$ provides a set of functions that can be used to assign a member of A to x. In this sense, we can think of $w \in H(U)$ as interpreting x in A, and of U as providing the context in which x is given this interpretation. If in addition we think of the members of $H(U)$ as states of affairs, then w represents a state of affairs where, given the context U, x is given a particular interpretation in A. In this sense, w indicates what x might mean in terms of A, given the context U.

In view of the work we shall need to complete in Chapter 3, given a state v of affairs, we shall be interested in those states of affairs that preserve the interpretation given to x by v in the context U. Informally, by this we mean that for $w \in H(U)$, $w(x)$ should be "close by" $v(x)$ (compare the discussion following Definition 3.21). To provide a formal notion of what it means to be "close by" within $\mathbf{A}$, let us presuppose a set-to-set function $C_A : \mathcal{P}(A) \longrightarrow \mathcal{P}(A)$ on $\mathbf{A}$. Typically, C_A will be a closure operator on A. We shall then be interested in those $w \in H(U)$ such that $w(x) \in C_A(\{v(x)\})$. Definition 2.30 may be compared to Definition 2.4 in [12] and also the identity (5) in [16].

Definition 2.30. For $x \in X$ and $S \subseteq X$, let $U \in \Omega X$ be such that $x \in U$ and $S \subseteq U$, and let $v, w \in H(U)$. We say that, given the context U,

i) w *satisfies* x *relative to* v if and only if $w(x) \in C_A(\{v(x)\})$, and we write $w \models_v x$ and call w a *model of* x *relative to* v. If x has a model relative to v, then x is *satisfiable relative to* v otherwise it is a *contradiction relative to* v. If every w in $H(U)$ satisfies x relative to v then x is called a *tautology relative to* v.

ii) w *satisfies* S *relative to* v if and only if $w(S) \subseteq C_A(\{v(S)\})$, and we write $w \models_v S$ and call w a *model of* S *relative to* v. We say that S is *consistent relative to* v if it has a model relative to v, otherwise it is *inconsistent relative to* v.

We refer to v as the *reference interpretation*.

As a notational convenience, we suppress the letter v in the subscripting and assume that satisfaction is determined relative to a given reference interpretation v. The context $U \in \Omega X$ will be allowed to range over the open sets of X, however.

For $U \in \Omega X$ and $x \in U$, let us define

$$W_{x,U} = \{w \in H(U) \mid w \models_v x\}$$

We shall think of $W_{x,U}$ as the meaning of x (relative to v) given the context U. By extension, for $S \subseteq U$ we define

$$W_{S,U} = \{w \in H(U) \mid w \models_v S\}$$

and we take $W_{S,U}$ to be the meaning of S (relative to v) given the context U.

We can order the members of X by comparing the values assigned to them by their meanings, independently of any context.

Definition 2.31. For $x, y \in X$ and $S, T \subseteq X$, we define $\leq_0 \subseteq X \times X$ to be such that

$$x \leq_0 y \quad \text{if and only if} \quad \forall U \in \Omega X \mid x, y \in U. \forall u \in W_{x,U} \cup W_{y,U}.[u(y) \in C_A(\{u(x)\})]$$

We define the *order of* A*-entrenchment* $\leq_e \subseteq X \times X$ to be the transitive closure of $\leq_0$. That is,

$$x \leq_e y \quad \text{if and only if} \quad \exists x_1, x_2, \ldots, x_n \in X.[x \leq_0 x_1 \leq_0 x_2 \leq_0 \ldots \leq_0 x_n \leq_0 y]$$

Correspondingly, we define $\leq_0^{\uparrow} \subseteq \mathcal{P}(X) \times \mathcal{P}(X)$ such that

$$S \leq_0^{\uparrow} T \quad \text{if and only if} \quad \forall U \in \Omega X \mid S, T \subseteq U. \forall u \in W_{S,U} \cup W_{T,U}.[u(T) \subseteq C_A(u(S))]$$

and we define $\leq_e^{\uparrow}$ as the transitive closure of $\leq_0^{\uparrow}$ so that

$$S \leq_e^{\uparrow} T \text{ if and only if } \exists X_1, X_2, \ldots, X_n \in \mathcal{P}(X).[S \leq_0^{\uparrow} X_1 \leq_0^{\uparrow} X_2 \leq_0^{\uparrow} \ldots \leq_0^{\uparrow} X_n \leq_0^{\uparrow} T]$$

As before, if $\mathbf{A}$ is understood, we shall refer to the order of A-entrenchment as just the order of entrenchment.

From Definition 2.31, $\leq_e$ is a quasi-order. In general, anti-symmetry is not present. We can recover a partial order from $\leq_e$ by defining an equivalence relation $\approx$ on X such that $x \approx y$ whenever $x \leq_e y$ and $y \leq_e x$. We shall, however, defer this aspect to Section 3.2.1 where we recover a partial order by identifying sentences under a notion of logical equivalence.

In the worked example of Section 4.2, we shall exhibit a second order on X that is also derived from the meanings of the members of X (see Definition 4.11). Consequently we shall denote the order on X generically as $\leq_X$, indicating separately the mechanism of its derivation.

We are interested in the special case where ΩX coincides with the Alexandroff topology on X with respect to $\leq_X$. The family of sets of the form $\uparrow x$, for any $x \in X$, is a base for this topology, and every open set of the topology is also an up-set with respect to $\leq_X$ (compare p101 in [3] and also p122 in [8]). Conveniently, each open set of the topology is then also a relational structure, since for any $U \in \Omega X$, $x \in U$ and $y \in X$, if $x \leq_X y$ then $y \in U$ as well. The significance of this feature will be apparent from the work we shall set out in Section 2.3.

Definition 2.32. Let (X, H) be an r-algebraic space over $\mathbf{A}$, and let $\leq_X$ be the order on X derived by Definition 2.31.

1. If every open set $U \in \Omega X$ is an up-set with respect to $\leq_X$, and every up-set with respect to $\leq_X$ is also an open set, then we call $(X, H, \leq_X)$ an *ordered space over* $\mathbf{A}$.

2. A morphism $f : (X, H_X, \leq_X) \longrightarrow (Y, H_Y, \leq_Y)$ of ordered spaces over $\mathbf{A}$ is given by a continuous function $k : X \longrightarrow Y$ that is also a homomorphism from $(X, \leq_X)$ to $(Y, \leq_Y)$, together with a k-morphism $g : H_Y \longrightarrow H_X$ of sheaves of $\mathbf{A}$-simulations. It is an *isomorphism of ordered spaces over* $\mathbf{A}$, and $(X, H_X, \leq_X)$ and $(Y, H_Y, \leq_Y)$ are *isomorphic as ordered spaces over* $\mathbf{A}$, if k is a homeomorphism and an isomorphism from $(X, \leq_X)$ to $(Y, \leq_Y)$, and $g : H_Y \longrightarrow H_X$ is an isomorphism of sheaves of $\mathbf{A}$-simulations.

If $\mathbf{A}$ is understood, we shall abbreviate "ordered space over $\mathbf{A}$" to "ordered space". By abuse of terminology we shall usually write "X is an ordered space (over $\mathbf{A}$)".

Now suppose that we are given a set X. For work we shall need to complete in Chapters 3 and 4, we require that X should have an algebraic structure of type $\mathcal{F}$. Starting from just the algebra (X, F) of type $\mathcal{F}$, we can use $\mathbf{A}$ to derive $\leq_X$, ΩX and the structure sheaf H by applying some modest specialisations, set out below, to the exposition of ordered spaces we have just recorded. The outcome is then, by Definition 2.32, an ordered space. Because X is equipped with an algebraic structure, $(X, F, \leq_X)$ is in fact an ordered r-algebra, and to emphasise this property we refer to $(X, H, \leq_X)$ as an ordered r-algebraic space.

We allow X to arise from a collection of primitive objects. Definition 2.33 is based on Definition 10.1 in Chapter II of [10]. We have used the words "primitive" and "sentence" where [10] has used "variable" and "term".

Definition 2.33. Let $\mathcal{F}$ be a language of algebras and let S_0 be a set of objects called *primitives*. A *sentence of type* $\mathcal{F}$ *over* S_0 is defined inductively as

- Every member of $S_0 \cup \mathcal{F}_0$ is a sentence of type $\mathcal{F}$.
- If $f \in \mathcal{F}$ is an n-ary function symbol, where $n \geq 1$, and $x_1, x_2, \ldots, x_n$ are sentences of type $\mathcal{F}$, then $f(x_1, x_2, \ldots, x_n)$ is a sentence of type $\mathcal{F}$.
- Nothing else is a sentence of type $\mathcal{F}$ except by virtue of the above.

The smallest set to contain all sentences of type $\mathcal{F}$ over S_0 is denoted by $\mathcal{S}_{\mathcal{F}}(S_0)$. For any sentence $s \in \mathcal{S}_{\mathcal{F}}(S_0)$, the *language of* s is the set L_s defined such that

- If $s \in S_0 \cup \mathcal{F}_0$ then $L_s = \{s\}$.
- If s has the form $f(x_1, x_2, \ldots, x_n)$, where f is an n-ary function symbol, where $n \geq 1$, and $x_1, x_2, \ldots, x_n$ are sentences of type $\mathcal{F}$, then $L_s = L_{x_1} \cup L_{x_2} \cup \ldots \cup L_{x_n}$.

For any non-empty set $U \subseteq \mathcal{S}_{\mathcal{F}}(S_0)$ of sentences of type $\mathcal{F}$, the language of U is $L_U = \bigcup_{x \in U} L_x$.

For our work, we shall take S_0 to be a non-empty, countable set of distinct primitives.

The set $\mathcal{S}_{\mathcal{F}}(S_0)$ of sentences of type $\mathcal{F}$ over S_0 gives rise to an algebra of type $\mathcal{F}$ in the expected manner. Definition 2.34 is based on Definition 10.4 in Chapter II of [10].

Definition 2.34. Let $\mathcal{F}$ be a language of algebras and let S_0 be a set of objects called *primitives*. The *sentence algebra* $\mathbf{S} = (S, F)$ *of type* $\mathcal{F}$ *over* S_0 has $S = \mathcal{S}_{\mathcal{F}}(S_0)$, and the fundamental operations of $\mathbf{S}$ satisfy

$$f^{\mathbf{S}} : S^n \longrightarrow S : (x_1, x_2, \ldots, x_n) \mapsto f(x_1, x_2, \ldots, x_n)$$

for $x_1, x_2, \ldots, x_n \in S$ and $f \in \mathcal{F}_n$ for $n \geq 1$. If f is a nullary function symbol corresponding to the sentence $a_f \in \mathcal{S}_{\mathcal{F}}(S_0)$, then $f^{\mathbf{S}} : \emptyset \longrightarrow \{a_f\}$.

The sentence algebra $\mathbf{S}$ of type $\mathcal{F}$ over S_0 may be thought of as being generated by S_0 (compare p71 in [10]). Furthermore, by Theorem 10.8 in [10], $\mathbf{S}$ is freely generated by S_0 (compare Definition 3.4 in [10]). The mechanism by which $\mathbf{S}$ is generated from S_0 is made formal in Section 3 of Chapter II of [10]. The procedure is not germane to our work, however, so we will not provide an exposition for it here.

From Theorem 6.2 in Chapter II of [10], we have the following. The proof of this result is supplied in [10], so we will not present it here.

Theorem 2.35. *Let $\mathbf{S}$ be an algebra of type $\mathcal{F}$ generated by a set S_0 and let $\mathbf{T}$ be an algebra of type $\mathcal{F}$. If $f, g \in [S \longrightarrow T]$ are homomorphisms from $\mathbf{S}$ to $\mathbf{T}$ such that for any $s \in S_0$, $f(s) = g(s)$, then $f = g$.*

Let $\mathbf{X} = (X, F)$ be the sentence algebra of type $\mathcal{F}$ over a non-empty, countable set X_0 of primitives. A homomorphism from $\mathbf{X}$ to $\mathbf{A}$ we shall call an *A-assignment (over X)*, and the restriction to X_0 of an A-assignment over X we shall call an *A-valuation (over X_0)*. If $\mathbf{A}$ is understood, we shall refer to A-assignments and A-valuations as just assignments and valuations. The family of valuations over X_0 is denoted by $[X_0 \longrightarrow A]$, and for any non-empty set $U \subseteq X_0$, we obtain the family $[U \longrightarrow A]$ of valuations over U by restriction.

From Theorem 2.35, each valuation over X_0 uniquely determines an assignment over X, so to determine the image of a sentence $x \in X$ we need only consider the action of the corresponding valuation on the language L_x of x. Given a valuation v and an arbitrary sentence $x \in X$, we shall abuse notation and write $v(x)$ to represent the value assigned to x by the assignment determined by v.

As a technical convenience, we shall regard the values in X that correspond to the nullary operations of $\mathcal{F}$ as primitives in X_0 as well. Because $X = \mathcal{S}_{\mathcal{F}}(X_0)$, for every subset U of X_0 there is $x \in X$ such that $L_x = U$. Consequently, we topologise X_0 with the discrete topology, and form the sheaf G on X_0 such that for any $U \in \Omega X_0$, $G(U) = [U \longrightarrow A]$, the valuations over U. For any $V \in \Omega X_0$ with $V \subseteq U$, we have the restriction map $\rho^U_V : G(U) \longrightarrow G(V) : s \mapsto s_{|V}$. Trivially, ρ^U_U is the identity map on U, and for any $W \in \Omega X_0$ with $W \subseteq V$, $\rho^U_W = \rho^V_W \circ \rho^U_V$. Furthermore, it follows readily that G satisfies the monopresheaf and glueing conditions of Definition 2.2, so G is in fact a sheaf, as intended. We shall refer to G as the sheaf of A-valuations over X_0. By Proposition 2.29 and Definition 2.27, (X_0, G) is an r-algebraic space over $\mathbf{A}$.

Suppose that we have a reference valuation $v \in [X_0 \longrightarrow A]$. For any $x \in X$, to determine the meaning of x in the setting of valuations, we must consider those valuations that can assign a value to x. That is, we consider those valuations in $[U \longrightarrow A]$ with $U \in \Omega X_0$ and $L_x \subseteq U$. Similarly, for $S \subseteq X$, we consider those valuations in $[U \longrightarrow A]$ with $U \in \Omega X_0$ and $L_S \subseteq U$. Intuitively, the sets $W_{x,U}$ and $W_{S,U}$ defined earlier now take the form

$$\begin{array}{rcll} W_{x,U} & = & \{w \in G(U) \mid w \models_v x\} & \text{for any } U \in \Omega X_0 \text{ with } L_x \subseteq U \\ W_{S,U} & = & \{w \in G(U) \mid w \models_v S\} & \text{for any } U \in \Omega X_0 \text{ with } L_S \subseteq U \end{array}$$

and for Definition 2.31, the condition on U is strengthened so that $W_{x,U}$ and $W_{y,U}$ are formed from any $U \in \Omega X_0$ with $L_x, L_y \subseteq U$, and similarly, $W_{S,U}$ and $W_{T,U}$ are formed for any $U \in \Omega X_0$ with $L_S, L_T \subseteq U$. The derivation of the order $\leq_X$ then proceeds identically. Given $\leq_X$, we then choose ΩX to be the Alexandroff topology with respect to $\leq_X$. Since X is countable, ΩX has a countable base and so $(X, \Omega X)$ is second-countable.

Next, let H be a map from the open sets of X to sets of valuations over subsets of X_0 such that for any $U \in \Omega X$, $H(U) = G(L_U)$. For $V \in \Omega X$ with $V \subseteq U$, let $\rho^U_V : H(U) \longrightarrow H(V) : s \mapsto s_{|L_V}$ as before. Then it is readily established that H is a presheaf of **A**-simulations over X, and also that H satisfies the monopresheaf and glueing conditions of Definition 2.2, so H is in fact a sheaf of **A**-simulations over X. Finally, as in Section 2.1.2, if we let $\mathcal{U}_x = \{U \in \Omega X \mid x \in U\}$ for $x \in X$, then the pair

$$\mathbf{U} = (\{H(U)\}_{U \in \mathcal{U}_x}, \{\rho^U_V \mid U, V \in \mathcal{U}_x \text{ and } V \subseteq U\})$$

is a direct system of sets with direct limit

$$\mathbf{V} = (G(L_x), \{\alpha_U : H(U) \longrightarrow G(L_x) : s \mapsto s_{L_x} \mid U \in \mathcal{U}_x\})$$

from which it follows that the stalk H_x of H at x is just $G(L_x)$. We shall refer to H as the sheaf of A-valuations over X. The pair (X, H) is then an r-algebraic space of type $\mathcal{F}$ over **A**. Together with the order $\leq_X$, by Definition 2.32, $(X, H, \leq_X)$ is an ordered space over **A**.

Definition 2.36. Let $\mathcal{F}$ be a language of algebras, and let **A** be a bounded r-algebra of type $\mathcal{F}$ such that $\leq_A$ is a partial order.

1. An *ordered r-algebraic space of type $\mathcal{F}$ over* **A** is an ordered space $(X, H, \leq_X)$, where $\leq_X$ is derived by Definition 2.31, such that $(X, F, \leq_X)$ is an ordered r-algebra of type $\mathcal{F}$.

2. A morphism $(X, H_X, \leq_X) \longrightarrow (Y, H_Y, \leq_Y)$ of ordered r-algebraic spaces of type $\mathcal{F}$ over **A** is a morphism of ordered spaces over **A** for which the continuous function is a homomorphism from (X, F) to (Y, F). It is an isomorphism if it is an isomorphism of ordered spaces such that the continuous function is an isomorphism from (X, F) to (Y, F).

We shall abbreviate the term "ordered r-algebraic space of type $\mathcal{F}$ over **A**" to "ordered r-algebraic space (over **A**)", omitting any reference to the type $\mathcal{F}$ and underlying algebra **A** if this causes no confusion. By abuse of terminology, we shall usually write "X is an ordered r-algebraic space (over **A**)".

2.3 Ordered Relational Algebraic Manifolds

In this section, we define a manifold in terms of a family of ordered spaces over a given bounded r-algebra **A**. The definition that we provide is an adaptation of the sheaf-theoretic definition of a manifold given in [52], and allows the manifold to inherit an order from the ordered spaces. We conclude by setting out a construction procedure for these manifolds.

2.3.1 Ordered Manifolds over a Relational Algebraic Structure

As in Section 2.2.2, we presuppose a bounded r-algebra $\mathbf{A} = (A, F, \mathbf{0}_A, \mathbf{1}_A, \leq_A)$ of type $\mathcal{F}$ in which $\leq_A$ is a partial order.

Informally, a local homeomorphism between two topological spaces X and Y ensures that locally, X is topologically and set-theoretically the same as Y. If X and Y are both ordered spaces, however, the structure sheaves need to be "locally the same" as well. To this end, local isomorphisms allow us to say when two ordered spaces are locally equivalent. Definition 2.37 may be compared to Definition 4.3.6 in [52].

Definition 2.37. Let $(M, G, \leq_M)$ be an ordered space over $\mathbf{A}$. An ordered space $(X, H, \leq_H)$ over $\mathbf{A}$ is *locally isomorphic* to $(M, G, \leq_M)$ if and only if for all $x \in X$ there is a neighbourhood $U \in \Omega X$ of x and an open set V in M such that $(U, H_{|U}, \leq_{X|U})$ and $(V, G_{|V}, \leq_{M|V})$ are isomorphic as ordered spaces over $\mathbf{A}$.

Remark 2.38. Definition 2.37 makes the importance of the Alexandroff topology clearer. The open sets U and V are both up-sets, and since $(U, H_{|U}, \leq_{X|U})$ and $(V, G_{|V}, \leq_{M|V})$ are isomorphic as ordered spaces over $\mathbf{A}$, U carries the same order as V.

We may now define an ordered manifold (compare Definition 4.3.6 in [52]).

Definition 2.39. Let M be a class of ordered spaces over $\mathbf{A}$. An ordered space $(X, H, \leq_X)$ over $\mathbf{A}$ is an *ordered manifold of type* M *over* $\mathbf{A}$ if and only if there is an open cover $\mathcal{U}$ of X such that for each $U \in \mathcal{U}$, $(U, H_{|U}, \leq_{X|U})$ is locally isomorphic to some $(M, G, \leq_M) \in \mathsf{M}$.

In view of Remark 2.38, we may see an ordered manifold of type M over $\mathbf{A}$ as inheriting the order $\leq_X$ from the ordered spaces in M.

As before, we abbreviate "ordered manifold of type M over $\mathbf{A}$" to "ordered manifold", omitting the references to M and $\mathbf{A}$ if these are understood and doing so causes no confusion. By abuse of terminology we shall usually write "X is an ordered manifold". A morphism (isomorphism) between ordered manifolds of type M over $\mathbf{A}$ may then be expressed in terms of morphisms (isomorphisms) of ordered spaces over $\mathbf{A}$.

Remark 2.40. Manifolds are usually presented as topological n-manifolds, which are T_2, second-countable topological spaces locally homeomorphic to $\mathbb{R}^n$ with the Euclidean topology. The T_2 condition is required to exclude certain pathological spaces such as the "line with two origins" (see Example 13.9(b) in [56], Exercise 3-8 in [34] and also Remark 4.3.8 in [52]), and also to ensure that certain kinds of analysis can be carried out on a manifold. The T_2 condition ensures that there are enough open sets to work with, while second-countability ensures that there are not too many. Definition 2.39 can be phrased in terms of what are called geometric spaces in [52]. Given a commutative ring $\mathbf{R}$ with identity, a geometric space over $\mathbf{R}$ is a topological space X together with a sheaf H of $\mathbf{R}$-algebras (compare Definition 2.27). An $\mathbf{R}$-algebra is a pair $(\mathbf{S}, f)$, where $\mathbf{S}$ is a commutative ring with identity and $f : \mathbf{R} \longrightarrow \mathbf{S}$ is a $\mathbf{1}$-preserving ring morphism (compare the $\mathbf{A}$-simulations of Definition 2.24). The sheaf H is such that the stalks are also local rings. If in Definition 2.39 M is instead a class of geometric spaces over $\mathbf{R}$, we recover the manifolds of Definition 4.3.6 in [52]. It is from these manifolds that a topological n-manifold (with, for example, a smooth structure) can be recovered (see Example 4.3.7 in [52]).

With the manifolds of Definition 2.39, we have been able to drop the requirement that a manifold should resemble $\mathbb{R}^n$ locally. As described in Section 1.3, the relaxation of this requirement was important to us because we would need to use topologies with a local structure different from that of $\mathbb{R}^n$ with the Euclidean topology.

Given an ordered manifold X of type M, we can use the ordered spaces in M to explore the structure of X. Let us presuppose a class M of ordered spaces over $\mathbf{A}$ and an ordered manifold $(X, H, \leq_X)$ of type M over $\mathbf{A}$. Definition 2.41 is based on pp4-5 and p12 of [33].

Definition 2.41. A *(coordinate) chart* on X is a pair (U, ϕ) where $U \in \Omega X$ and $\phi : U \longrightarrow V$, where $V = \phi(U)$ is an open subset of some $(M, G, \leq_M)$ in M, such that $(U, H_{|U}, \leq_{X|U})$ and $(V, G_{|V}, \leq_{M|V})$ are isomorphic as ordered spaces over $\mathbf{A}$. The set U is called the *(coordinate) domain* of the chart, and the map ϕ is called a *(local) coordinate map*. If (U, ϕ) is a chart on X and $x \in U$, we say that (U, ϕ) *covers* x. An *atlas for* X is a collection of charts on X whose domains cover X. That is, if $\mathcal{A} = \{(U_i, \phi_i)\}_{i \in I}$ is an atlas for X, then $\{U_i\}_{i \in I}$ is an open cover of X.

A chart (U, ϕ) on X provides a view of X, since for the region U, we know that X is locally the same as $\phi(V)$, which is an open set in some $M \in$ M. If two charts have overlapping domains, we can then translate between these views. Definition 2.42 is based on p12 in [33].

Definition 2.42. Let (U, ϕ) and (V, ψ) be any two charts on X such that $U \cap V \neq \emptyset$. The maps

$$\begin{array}{ll} \psi \circ \phi^{-1} : \phi(U \cap V) \longrightarrow \psi(U \cap V) & \text{and} \\ \phi \circ \psi^{-1} : \psi(U \cap V) \longrightarrow \phi(U \cap V) & \end{array}$$

are called *transition maps*.

Remark 2.43. As a notational convenience we shall denote charts with lower case letters such as c and d. We write U_c and ϕ_c for the domain and coordinate map of a chart c. We write c_U to mean the ordered space $(U_c, H_{|U_c}, \leq_{X|U_c})$, and if $V = \phi_c(U_c)$ is an open set in $(M, G, \leq_M) \in$ M, we write c_ϕ to mean the ordered space $(V, G_{|V}, \leq_{M|V})$. Given two charts c and d, if $U_c \cap U_d \neq \emptyset$ we write $c \circ d^{-1}$ and $d \circ c^{-1}$ to mean the charts $(U_c \cap U_d, \phi_c \circ \phi_d^{-1})$ and $(U_c \cap U_d, \phi_d \circ \phi_c^{-1})$ respectively. Finally, given a chart c, if V is an open subset of U_c, we write $c_{|V}$ to mean $(V, (\phi_c)_{|V})$, the restriction of c to V.

Where two charts overlap, our exploration of X should be independent of the choice of chart. This will be the case if the charts are compatible in the following sense (compare p12 in [33]).

Definition 2.44. Let c and d be two charts in an atlas $\mathcal{A}$ for X. Then c and d are *order-compatible* if either $U_c \cap U_d = \emptyset$ or the transition maps $\phi_d \circ \phi_c^{-1}$ and $\phi_c \circ \phi_d^{-1}$ are such that $(c \circ d^{-1})_\phi$ and $(d \circ c^{-1})_\phi$ are isomorphic as ordered spaces over $\mathbf{A}$.

Order-compatibility between charts may naturally be extended to order-compatibility between atlases (compare pp13–14 in [33]).

Definition 2.45. Let $\mathcal{A}$ be an atlas for X. If any two charts in $\mathcal{A}$ are order-compatible, we call $\mathcal{A}$ an *order-compatible atlas for* X. If $\mathcal{A}$ and $\mathcal{B}$ are two order-compatible atlases for X, then $\mathcal{A}$ and $\mathcal{B}$ are *jointly order-compatible* if and only if $\mathcal{A} \cup \mathcal{B}$ is an order-compatible atlas for X. An order-compatible atlas is a *maximal order-compatible atlas* if it is not properly contained in any other order-compatible atlas. Given an order-compatible atlas $\mathcal{A}$, the *maximal order-compatible atlas of* $\mathcal{A}$ is an atlas that contains all possible charts on X that are order-compatible with the charts in $\mathcal{A}$.

A given atlas may not determine an ordered manifold uniquely. As an example, $\{(X, \mathrm{id}_X)\}$ and $\{(\uparrow x, \mathrm{id}_{\uparrow x}) \mid x \in X\}$ are both atlases for X. We might therefore expect to specify a maximal order-compatible atlas for X. The next result shows that we need only specify *some* order-compatible atlas, as this will determine a unique maximal order-compatible atlas for X (compare Lemma 1.10 on p14 in [33]).

Proposition 2.46. *Every order-compatible atlas for X is contained in a unique maximal order-compatible atlas for X. Two order-compatible atlases for X determine the same maximal order-compatible atlas for X if and only if they are jointly order-compatible.*

Proof. Let $\mathcal{A}$ be an order-compatible atlas for X, and let $\mathcal{A}'$ denote the set of all charts on X that are order-compatible with every chart in $\mathcal{A}$. Let c and d be any two charts in $\mathcal{A}'$. If $U_c \cap U_d = \emptyset$, then c and d are trivially compatible by Definition 2.44. Otherwise, for any $x \in U_c \cap U_d$, let b be a chart in $\mathcal{A}$ such that $x \in U_b$. Then also $V = U_c \cap U_d \cap U_b \neq \emptyset$. Because c and d are order-compatible with b, it follows that $(c_{|V})_\phi$, $(d_{|V})_\phi$ and $(b_{|V})_\phi$ are all isomorphic as ordered spaces. Since the choice of x was arbitrary, by Definition 2.37 it follows that $(c_{|U_c \cap U_d})_\phi$ and $(d_{|U_c \cap U_d})_\phi$ are isomorphic as ordered spaces, and hence c and d are order-compatible and $\mathcal{A}'$ is an order-compatible atlas. It is maximal since any chart that is order-compatible with every chart in $\mathcal{A}$ is contained in $\mathcal{A}'$. It is unique because if $\mathcal{B}$ is another maximal order-compatible atlas, then every chart in $\mathcal{A}'$ is in $\mathcal{B}$ and every chart in $\mathcal{B}$ is in $\mathcal{A}'$, whence $\mathcal{A}' = \mathcal{B}$. For the second part, let $\mathcal{A}$ and $\mathcal{B}$ be two order-compatible atlases for X. For the forward direction, let $\mathcal{C}$ be the unique maximal order-compatible atlas determined by $\mathcal{A}$ and $\mathcal{B}$. Then $\mathcal{A} \subseteq \mathcal{C}$ and $\mathcal{B} \subseteq \mathcal{C}$, so $\mathcal{A} \cup \mathcal{B} \subseteq \mathcal{C}$. Since $\mathcal{C}$ is an order-compatible atlas, any two charts in $\mathcal{A} \cup \mathcal{B}$ must then be order-compatible, so $\mathcal{A} \cup \mathcal{B}$ is an order-compatible atlas, and $\mathcal{A}$ and $\mathcal{B}$ are jointly order-compatible. For the reverse direction, let $\mathcal{A}'$ and $\mathcal{B}'$ be the maximal atlases of $\mathcal{A}$ and $\mathcal{B}$. Then since $\mathcal{A}$ and $\mathcal{B}$ are jointly order-compatible, $\mathcal{B} \subseteq \mathcal{A}'$ and $\mathcal{A} \subseteq \mathcal{B}'$. But since $\mathcal{A}'$ and $\mathcal{B}'$ are order-compatible atlases for X, every chart in $\mathcal{A}'$ is order-compatible with every chart in $\mathcal{B}$, so that $\mathcal{A}' \subseteq \mathcal{B}'$. Similarly $\mathcal{B}' \subseteq \mathcal{A}'$, and hence $\mathcal{A}' = \mathcal{B}'$. □

Suppose now that instead of $(X, H, \leq_X)$, we are given only the set X. We would like to topologise X and equip it with an order $\leq_X$ and a structure sheaf H of $\mathbf{A}$-simulations in such a way that $(X, H, \leq_X)$ is an ordered manifold of type M. In the next section, we describe how this may be accomplished.

2.3.2 Construction of an Ordered Relational Algebraic Manifold

To endow a given set X with the structure of an ordered manifold, we follow the same procedure as set out in Lemma 1.23 on pp21-22 of [33]. This lemma was developed for the case of a topological n-manifold with the aim of conferring on it a smooth structure. We have adapted it to match the ordered manifolds of Definition 2.39.

For convenience, we shall write the class M of ordered spaces as $\{(M_j, G_j, \leq_j)\}_{j\in J}$. It is possible that not all of the $\leq_j$ will have arisen from the same bounded r-algebra. Assume then that each M_j is an ordered space over the distinct bounded r-algebra $\mathbf{A}_j$. Each G_j is then a sheaf of $\mathbf{A}_j$-simulations.

We let $\mathbf{A}$ be a bounded r-algebra such that each $\mathbf{A}_j$ can be embedded in $\mathbf{A}$. Observe that, by the conventions established after Definition 2.21, whereby nullary functions are preserved by homomorphisms, the top and bottom elements of each $\mathbf{A}_j$ are mapped to the top and bottom elements of $\mathbf{A}$. From the discussion preceding Definition 2.25, each G_j can then be expressed as a sheaf of $\mathbf{A}$-simulations, which gives us a class of ordered spaces over $\mathbf{A}$. We shall continue to use the notation $\mathsf{M} = \{(M_j, G_j, \leq_j)\}_{j\in J}$ for this class and its members.

Next, let $\mathcal{U} = \{U_i\}_{i\in I}$ be a family of subsets of X such that $X = \bigcup_{i\in I} U_i$.

Remark 2.47. If we require that X, once topologised, should be a second-countable space (*i.e.* X should have a countable base), $\mathcal{U}$ should be such that countably many of the U_i cover X, and each member of M must be second-countable also. The M_j are second-countable if each arises as the sentence algebra of type $\mathcal{F}$ over a countable set of primitives (see the discussion following Theorem

2.35). If we require the T_2 condition to be met, each M_j must be T_2 and $\mathcal{U}$ must satisfy the additional condition that for any distinct $x, y \in X$, either there is $U \in \mathcal{U}$ with $x, y \in U$ or there are disjoint $U, V \in \mathcal{U}$ with $x \in U$ and $y \in V$. Both properties of ΩX are then readily established, and we will not impose either of these requirements here.

Fix a collection $\{\phi_i\}_{i\in I}$ of injective maps such that

M1: for each $i \in I$, $\phi_i : U_i \longrightarrow M_{j_i}$, where $j_i \in J$, and $\phi_i(U_i) \in \Omega M_{j_i}$

M2: for each $i, k \in I$, $\phi_i(U_i \cap U_k) \in \Omega M_{j_i}$ and $\phi_k(U_i \cap U_k) \in \Omega M_{j_k}$, where $j_i, j_k \in J$

M3: for each $i, k \in I$, if $U_i \cap U_k \neq \emptyset$ then $(V_i, (G_{j_i})_{|V_i}, (\leq_{j_i})_{|V_i})$ and $(V_k, (G_{j_k})_{|V_k}, (\leq_{j_k})_{|V_k})$, where $V_i = \phi_i(U_i \cap U_k)$ and $V_k = \phi_k(U_i \cap U_k)$, are isomorphic as ordered spaces

Note that from **M3**, for each $i, k \in I$, if $U_i \cap U_k \neq \emptyset$ then the transition maps

$$\begin{aligned} &\phi_i \circ \phi_k^{-1} : \phi_k(U_i \cap U_k) \longrightarrow \phi_i(U_i \cap U_k) \qquad \text{and} \\ &\phi_k \circ \phi_i^{-1} : \phi_i(U_i \cap U_k) \longrightarrow \phi_k(U_i \cap U_k) \end{aligned}$$

are homeomorphisms.

Now let $\mathcal{B}$ be the family of sets given by

$$\mathcal{B} = \bigcup_{i\in I} \{\phi_i^{-1}(V \cap \phi_i(U_i)) \mid V \in \Omega M_{j_i}, j_i \in J\}$$

Lemma 2.48. *The family $\mathcal{B}$ is a base for a topology on X.*

Proof. For any $i \in I$, $\phi_i(U_i) \in \Omega M_{j_i}$, and hence $U_i \in \mathcal{B}$. It follows that $\mathcal{U} \subseteq \mathcal{B}$, and since $X = \bigcup_{i\in I} U_i$, we obtain $X \subseteq \bigcup_{B\in\mathcal{B}} B$. For any $B \in \mathcal{B}$ we have $B = \phi_i^{-1}(V \cap \phi_i(U_i))$ for some $i \in I$, where $V \in \Omega M_{j_i}$. Since ϕ_i is injective and $V \cap \phi_i(U_i) \subseteq \phi_i(U_i)$, it follows that $B \subseteq U_i$. Consequently, for any $B \in \mathcal{B}$ there is $i \in I$ such that $B \subseteq U_i$, and hence $\bigcup_{B\in\mathcal{B}} B \subseteq \bigcup_{i\in I} U_i = X$. Taken together, we have $X = \bigcup_{B\in\mathcal{B}} B$. Next, consider $B_1, B_2 \in \mathcal{B}$ with $B_1 = \phi_i^{-1}(V_1 \cap \phi_i(U_i))$ and $B_2 = \phi_k^{-1}(V_2 \cap \phi_k(U_k))$ for $i, k \in I$, $V_1 \in \Omega M_{j_i}$, $V_2 \in \Omega M_{j_k}$ and $j_i, j_k \in J$. Suppose that $x \in B_1 \cap B_2$. Then since $U_i \cap U_k \neq \emptyset$, by **M2** we have $V_i = \phi_i(U_i \cap U_k) \in \Omega M_{j_i}$ and $V_k = \phi_k(U_i \cap U_k) \in \Omega M_{j_k}$. Then $V_k \cap V_2 \in \Omega M_{j_k}$, and by **M3**, $V_3 = \phi_i \circ \phi_k^{-1}(V_k \cap V_2 \cap \phi_k(U_k)) \in \Omega M_{j_i}$. Since $V_3 \cap V_i \cap V_1 \in \Omega M_{j_i}$, we then have $B_3 = \phi_i^{-1}(V_3 \cap V_i \cap V_1 \cap \phi_i(U_i)) \in \mathcal{B}$ with $x \in B_3$. It follows readily that $B_3 \subseteq B_1 \cap B_2$, and hence by Theorem 5.3 in [56] $\mathcal{B}$ is a base for a topology on X. □

From Theorem 7.9 in [56], we have the following result.

Lemma 2.49. *Let X and Y be topological spaces, and let $f : X \longrightarrow Y$ be a bijective function. Then the following are equivalent:*

i) The function f is a homeomorphism.

ii) For any $G \subseteq X$, $f(G) \in \Omega Y$ if and only if $G \in \Omega X$.

Proof. i)⇒ii). Because f is a homeomorphism, f is bijective and continuous and f^{-1} is also continuous. For $G \subseteq X$, assume that $f(G) \in \Omega Y$. Since f is continuous, $f^{-1}(f(G)) \in \Omega X$. Since f is bijective, $f^{-1}(f(G)) = G$, and hence $G \in \Omega X$. Similarly, if $G \in \Omega X$, then since f^{-1} is continuous,

$(f^{-1})^{-1}(G) \in \Omega Y$. Since f is bijective, $(f^{-1})^{-1}(G) = f(G)$, and hence $f(G) \in \Omega Y$.

ii)⇒i). For any $G \subseteq X$, suppose that $f(G) \in \Omega Y$ if and only if $G \in \Omega X$. We are given that f is bijective. For any $V \in \Omega Y$, taking $G = f^{-1}(V)$ we have $f(G) = f(f^{-1}(V)) = V \in \Omega Y$, and so $G \in \Omega X$ by hypothesis. For any $G \in \Omega X$ we have $(f^{-1})^{-1}(G) = f(G) \in \Omega Y$ by hypothesis. Thus, f is bijective and continuous and f^{-1} is also continuou, and it follows that f is a homeomorphism. □

Let ΩX denote the topology on X that arises from the base $\mathcal{B}$.

Lemma 2.50. *For each $i \in I$, ϕ_i is a homeomorphism.*

Proof. Let $V = \phi_i(U_i) \in \Omega M_{j_i}$, and let $(U_i, \Omega U_i)$ and $(V, \Omega V)$ be topological spaces for which ΩU_i and ΩV are the relative topologies inherited from X and M_{j_i}. From the definition of $\mathcal{B}$, it follows that for $G \subseteq U_i$, if $\phi_i(G) \in \Omega M_{j_i}$ then $G = \phi_i^{-1}(\phi_i(G)) \in \Omega U_i$. Conversely, if $G \in \Omega U_i$ then there is $\mathcal{G} \subseteq \mathcal{B}$ such that $G = \bigcup_{U \in \mathcal{G}} U$. But then $\phi_i(G) = \bigcup_{U \in \mathcal{G}} \phi_i(U) \in \Omega M_{j_i}$ because for each $U \in \mathcal{G}$, $\phi_i(U) \in \Omega M_{j_i}$. By Lemma 2.49, ϕ_i is then a homeomorphism from U_i to $\phi_i(U_i)$. □

For any $U_i \in \mathcal{U}$, we define $\leq_{U_i} \subseteq U_i \times U_i$ such that for any $x, y \in U_i$,

$$x \leq_{U_i} y \quad \text{if and only if} \quad \phi_i(x) \leq_i \phi_i(y)$$

By **M3**, if $U_i \cap U_j \neq \emptyset$ for $i, j \in I$, then

$$\phi_i(x) \leq_i \phi_i(y) \quad \text{if and only if} \quad \phi_j(x) \leq_j \phi_j(y)$$

so that $\leq_{U_i}$ and $\leq_{U_j}$ are compatible on $U_i \cap U_j$. In this way, an order is conferred on X. We shall denote this order by $\leq_X$, as before. Because the order on each $\phi_i(U_i)$ is derived by application of Definition 2.31, it follows that $\leq_X$ is derived in this manner as well.

Because each member of M is an ordered space, from Definition 2.32, each member of M has the Alexandroff topology. Based on $\leq_X$, because each member of M carries the Alexandroff topology we can deduce that every basic open set in $\mathcal{B}$ is an up-set. Hence every open set in ΩX is an up-set as well. From the definition of the ϕ_i and of $\mathcal{B}$, for any $x \in X$, $\uparrow x \in \mathcal{B}$ as well. For any up-set $U \subseteq X$ we may write $U = \bigcup_{x \in U} \uparrow x$, which, as an arbitrary union of open sets, is also open in X. Hence X also has the Alexandroff topology. This gives us the following result.

Lemma 2.51. *The topology ΩX is the Alexandroff topology on X.*

We now need to provide a sheaf of **A**-simulations on X. By Definition 2.39, any sheaf H such that $(U_i, H_{|U_i}, \leq_{X|U_i})$ is locally isomorphic to $(\phi_i(U_i), (G_{j_i})_{|\phi_i(U_i)}, \leq_{j_i|\phi(U_i)})$, where $U_i \in \mathcal{U}$, will do. To develop such a sheaf, we shall use our putative atlas $\mathcal{A} = \{(U_i, \phi_i)\}_{i \in I}$. Thus, let H be a presheaf of **A**-simulations on X.

For any $U \in \Omega X$, let $I_U = \{i \in I \mid U_i \cap U \neq \emptyset\}$, and let $\mathcal{U}_U = \{U_i\}_{i \in I_U}$. For any $U_i \in \mathcal{U}_U$, we shall require that

$$\begin{aligned} H_{|U_i \cap U} &= (\phi_i^{-1})_*(G_{j_i})_{|\phi_i(U_i \cap U)} \\ \rho^{U_i \cap U}_{U_i \cap V} &= (\rho_{j_i})^{\phi_i(U_i \cap U)}_{\phi_i(U_i \cap V)} \quad \text{where } V \in \Omega X \text{ and } V \subseteq U \end{aligned}$$

A given section of $H(U)$ is then described by local sections that are provided by sheaves that are accessed via the charts. That is,

$$\forall s \in H(U). \forall U_i \in \mathcal{U}_U. \exists s_i \in H(U_i \cap U)[\rho^U_{U_i \cap U}(s) = s_i]$$

In addition, if these local sections can be fitted together, they should give rise to a section in $H(U)$. That is, if $\{s_i\}_{i \in I_U}$ is a family of sections such that for each $i \in I_U$, $s_i \in H(U_i \cap U)$, then if

$$\rho^{U_i \cap U}_{U_i \cap U_j \cap U}(s_i) = \rho^{U_j \cap U}_{U_i \cap U_j \cap U}(s_j)$$

for any $U_i, U_j \in \mathcal{U}_U$ with $U_i \cap U_j \cap U \neq \emptyset$, then there is a section $s \in H(U)$ such that for all $i \in I_u$,

$$\rho^{U}_{U_i \cap U}(s) = s_i$$

Finally, we shall identify two sections of $H(U)$ if they can be described by the same family of local sections. That is,

$$\forall s, t \in H(U).[\forall i \in I_U.[\rho^{U}_{U_i \cap U}(s) = \rho^{U}_{U_i \cap U}(t)] \Leftrightarrow s = t]$$

(where we have written the above condition as an equivalence because the reverse implication is always true).

It follows that, for the cover by members of $\mathcal{U}_U$, H satisfies the glueing and monopresheaf conditions (Definition 2.2) on U. However, because the members of $\mathcal{U}$ are all open in ΩX, any open cover of U can be expressed as an open cover by members of $\mathcal{U}$. Since the choice of $U \in \Omega X$ was arbitrary, it follows that H satisfies the monopresheaf and glueing conditions for any open set, and so is a sheaf of **A**-simulations on X.

Trivially, for any $U_i \in \mathcal{U}$, U_i is a cover of itself, so the requirement that $H_{|U_i} = (\phi_i^{-1})_*(G_{j_i})_{|\phi_i(U_i)}$ means that $(U_i, H_{|U_i}, \leq_i)$ is locally isomorphic to $(\phi_i(U_i), (G_{j_i})_{|\phi_i(U_i)}, \leq_{j_i|\phi_i(U_i)})$, and since the U_i provide an open cover of X, H meets the requirements of Definition 2.39.

By Definition 2.32, $(X, H, \leq_X)$ is then an ordered space over **A**, and by Definition 2.39, it is an ordered manifold of type M over **A**. We state this as the following result, thereby summarising the outcome of the procedure we have just described.

Theorem 2.52. *The triple $(X, H, \leq_X)$ is an ordered manifold of type* M *over* **A**.

Corollary 2.53. *Each member of $\mathcal{A}$ is a chart on X, and the family $\mathcal{A}$ of charts is an order-compatible atlas for X.*

Proof. This result follows from Definitions 2.41, 2.44 and 2.45, property **M3** and the definition of H just given. □

Summary

In this chapter, we constructed the framework that we will use to model belief revision as a local, relational operation. We presented the framework as a manifold that was constructed from a family of topological spaces, each equipped with a quasi-order. The manifold was constructed in such a way that it inherited a quasi-order from the topological spaces in the given family. Our approach to the construction of the manifold appealed to the sheaf-theoretic definition of a manifold given in [52] and consequently relied heavily on techniques and constructions from the field of sheaf theory.

Relational algebraic structures, or r-algebras, were the primary component of our framework. These structures combined the idea that a set could carry a specified algebraic structure as well as a binary relation. We illustrated our idea by showing how lattices could be considered as r-algebras. Certainly, for simple structures like these, it was mostly straightforward to prescribe a binary relation. For more complex algebraic structures, however, the relation between the points would be less obvious.

R-algebras thus provided us with a convenient means of simultaneously transporting an algebraic and relational structure on a set.

A related notion was the idea that one r-algebra could simulate another. This idea allowed us to construct a topological space over a particular r-algebra, where the r-algebraic structure on the space was provided by a sheaf of these r-algebra simulations. We called these spaces r-algebraic spaces. We could certainly have equipped the given space with a sheaf of r-algebras to give it an r-algebraic structure. However, our purpose in using a sheaf of r-algebra simulations was to equip the space with the structure of a specified r-algebra.

By considering an r-algebra for which the relation was a partial order, we were then able to derive a quasi-order and a topology on a given set of points, and in so doing, to construct an r-algebraic space that was itself equipped with an order. We called these spaces ordered r-algebraic spaces, and it was from these spaces that we were able to construct the required manifold.

By using the carrier set of a sentence algebra over a given set of primitives, we could derive the quasi-order from a notion of valuation and satisfaction. In a typical logic, the notion of satisfaction is derived from the assignment of a designated value such as `true` to a given sentence by a particular valuation. In our case, we did not use such especially designated values, but appealed instead to the idea of "preservation of degrees of truth" as described in [16]. Consequently, in Definition 2.30 a sentence could be considered satisfied by a valuation in a particular context if the valuation preserved the value assigned to the sentence in the same context by a reference valuation.

From the notion of satisfaction, we developed the idea of the meaning of a sentence, and from these meanings we were able to derive an order for the sentences in our sentence algebra. With an order in place, we could then derive a topology for the sentence algebra. To this end, we selected the Alexandroff topology with the construction of the manifold in view, because with the Alexandroff topology all open sets are also up-sets, which in turn would allow the manifold to inherit an order more easily.

Chapter 3

Disposition and Logic

We now commence specialisation of the framework of Chapter 2. Our goal in this chapter is to develop a model of a subject as a collection of sentences equipped with a logic.

The family of sentences that we consider is developed as the algebra of sentences of a specified type over a countable set of primitives. To provide an order for these sentences, we appeal to the notion of disposition, or propositional attitude, as described in Section 1.1, which we are able to represent as a bounded r-algebra. By applying the work of Chapter 2 (in particular Section 2.2.2), we are able to induce an order on sentences with this bounded r-algebra. The induced order then represents a form of triage with respect to the given disposition.

We next develop a logic that interacts with the induced order on the sentences via the idea of "preservation of degrees of truth" as described in [16]. This idea allows us to formulate a logic in which the degrees of disposition assigned to the members of a set of sentences are, in a certain sense, preserved under logical consequence. Logical consequence imposes additional structure on the sentences, and from this structure we are then able to formulate conjunction, disjunction, negation and contradiction in terms of the fundamental operations of the algebra. We proceed as follows.

Chapter Guide:

Section 3.1: A Model of Disposition. A disposition represents a propositional attitude towards a proposition or state of affairs in which that proposition is true. Intuitively, a given disposition can be held more strongly towards some propositions than others, which allows us to order propositions. In this section, we develop a model of a disposition as a bounded r-algebra of type $\mathcal{F}$, which allows us to order the sentences of a given algebra of type $\mathcal{F}$. This algebra, together with the order is then a bounded r-algebra, which we use as the foundation for our model of a subject.

Section 3.2: A Logic of Disposition. In this section, we develop a logic for a subject. We begin by describing the properties of the logic that are required for us to develop a belief revision relation in the style of [36]. By appeal to the idea of "preservation of degrees of truth", we then develop a logic in such a way that the fundamental operations of the subject also serve as logical connectives.

3.1 A Model of Disposition

In Chapter 1, we illustrated a propositional attitude with an example where James was hopeful that there would be enough food for his guests. The example suggested a binary characterisation of the disposition "hopefulness", so that James either hoped or did not hope that there was enough food for everyone. We could not express how fervently he might have hoped that there was enough food, or whether he was more or less hopeful that there was enough food than that his guests would enjoy the meal he had prepared.

With a set such as $[0, 1] \subseteq \mathbb{R}$, we can express such degrees of dispostion with infinite refinement. However, this is more than we need and masks the qualitative character of the problem. What we want is a set of possible degrees of disposition that retains the simple qualitative nature of the two-valued case yet allows us to express more refined degrees of disposition.

We begin our development of a disposition by specifying the set of sentences with which we shall associate degrees of disposition. With a view to the development in Section 3.2, we specify that $\mathcal{F} = \{+, *, '\} = \{2, 2, 1\}$, and we let $\mathbf{X} = (X, +, *, ')$ be the sentence algebra of type $\mathcal{F}$ over a non-empty, countable set X_0 of primitives. In $\mathbf{X}$ we take $+$ and $*$ to be associative and commutative, and we treat $'$ as a form of inverse so that for any $x \in X$, $x'' = x$. To avoid ambiguity in any sentence, we assume that $'$ binds more tightly than $*$, which binds more tightly than $+$.

Suppose we have a set D, whose members we wish to assign to the sentences in X to indicate the extent to which an agent holds a given disposition towards a particular sentence (x, say). Such assignment presumes that we can know exactly the degree to which the agent holds the disposition towards x. It would make our model brittle, for any discrepancies between actual and assigned degrees of disposition would be amplified by computations we would need to perform based on the structure of x. We would also need to commit to a type for the members of D – for example, qualitative values such as hot or lukewarm – which is certain to attract controversy.

We can make our model more resilient and obviate the need to fix the type of the members of D by introducing an element of non-determinism (compare p62 in [8] for example). To a given sentence we now assign a subset of D, thereby assigning a degree of disposition based on a property that the degree should satisfy, rather than assigning the degree purely by value. We can then exploit an important duality between sets and members of sets. From [8] (see p12), the idea of this duality is that

i) any particular thing is determined by its properties (Leibniz' 'Principle of the Identity of Indiscernables')

ii) any property is determined by the set of all things having that property (the 'Extensionality Principle')

A subset $U \subseteq D$ thus represents a property. Correspondingly, given $d \in D$, the properties of d are just those subsets of D that contain d. The subset U can also be seen as an observation about D, and d satisfies U if $d \in U$. By assigning the observation U rather than the degree d to a given sentence x, we indicate that, no matter what degree of disposition the agent may *actually* hold towards x, that degree will satisfy the observation U.

There are several advantages to this approach. First, a richer characterisation is possible because there are more properties than degrees of disposition. Second, because our model is formulated in terms of properties rather than actual degrees of disposition, by analogy with [22] degrees are not known by value but by observations that they satisfy. In a sense, the use of properties allows us to work with dispositions in a mostly point-free setting (compare [54] and [55], for example), and obviates a commitment to a type for the members of D. Finally, degrees of disposition are not necessarily equipped with an algebraic structure, so it may not be possible, for example, to determine the degree of disposition assigned to the sentence $x * y$ from the degrees of disposition assigned to the sentences x and y. In contrast, properties have an inherent logical nature that readily facilitates such combination and manipulation, and in this sense they behave more like truth values than do degrees of disposition.

Informally, to develop our model of a disposition, we begin with a set D of values about which we assume nothing and know very little apart from a given set $\mathcal{D}_0$ of observations about the members of D. To ensure that we have enough observations to work with, we impose certain restrictions on $\mathcal{D}_0$. In addition, to ensure that we can relate an arbitrary observation about D to the given observations, we extend $\mathcal{D}_0$ to derive a family $\mathcal{D}$ of sets called a topped intersection structure on D. By appeal to the duality between sets and properties, we are then able to order the members of D in a way that has a pleasing interaction with the observations in $\mathcal{D}$.

By applying a particular closure operator to $\mathcal{D}$, we are able to extract certain sets of properties, or specifications, as the closed sets of the operator. We then us these "closed" specifications to develop a set of surrogates for degrees of disposition. Our exposition of these surrogates uses ideas set out in [9] and [15] as a starting point. Consequently, a surrogate is a combination of two specifications, and it is the interplay between these two specifications that characterises the degree of disposition represented by the surrogate. Finally, we show that, when ordered in a certain way, the set of surrogates is a complete lattice.

3.1.1 Degrees of Disposition, Observations, Order and Specifications

Let D be a non-empty set. To avoid a lack of differentiation amongst the degrees of disposition, we assume without further comment that D has more than one point. A family of subsets of D represents a collection of observations about D, and to avoid certain kinds of pathological cases (compare Remark 2.40), this family must meet certain requirements.

Definition 3.1. Let X be a set. A non-empty family $\mathcal{X} = \{U_i\}_{i \in I}$ of subsets of X is an *admissible family of observations* about X if

i) $X \in \mathcal{X}$

ii) for any two distinct points in X there is a set in $\mathcal{X}$ that contains one but not the other

iii) it contains at most countably many subsets of X

The observation X is called the *trivial observation*.

Let $\mathcal{D}_0$ be an admissible family of observations about D. We shall require the the existence of two special elements $d_{\mathbf{0}}$ and $d_{\mathbf{1}}$ in D, for which

$d_{\mathbf{0}} \triangleq$ the *trivial degree*, which satisfies only the trivial observation D, and
$d_{\mathbf{1}} \triangleq$ the *universal degree*, which satisfies every observation in $\mathcal{D}_0$

In general, these elements may not be present in D, and to ensure their existence we artificially affix them to D if this is necessary. For the rest of our exposition, we therefore assume that $d_{\mathbf{0}}, d_{\mathbf{1}} \in D$.

Given an arbitrary observation $V \subseteq D$, we would like to relate V to the observations in $\mathcal{D}_0$. For any $d \in D$, if d satisfies V then d also satisfies $U \in \mathcal{D}_0$ if $V \subseteq U$. The closest observation to V is then just $\bigcap\{U \in \mathcal{D}_0 \mid V \subseteq U\}$, and to ensure that $\mathcal{D}_0$ contains this observation, we close $\mathcal{D}_0$ under non-empty intersections to form the topped intersection structure $\mathcal{D}$ on D. Observe that the empty intersection is included in $\mathcal{D}$ because by Definition 3.1(i), $D \in \mathcal{D}$, so in effect $\mathcal{D}$ is closed under arbitrary intersections.

From Definition 2.33 and Corollary 2.32 in [13] we have the following.

Definition 3.2. Let X be a set and let $\mathcal{X}$ be a family of subsets of X, ordered by inclusion. If, for every non-empty family $\{X_i\}_{i\in I} \subseteq \mathcal{X}$ we have $\bigcap_{i\in I} X_i \in \mathcal{X}$, then $\mathcal{X}$ is called an *intersection structure* (or *$\bigcap$-structure*) on X. If $X \in \mathcal{X}$, then $\mathcal{X}$ is called *topped $\bigcap$-structure* on X.

It is straightforward to show that a topped $\bigcap$-structure $\mathcal{X}$ on X is a complete lattice in which

$$\begin{aligned}\bigwedge_{i\in I} X_i &= \bigcap_{i\in I} X_i \\ \bigvee_{i\in I} X_i &= \bigcap\{Y \in \mathcal{X} \mid \bigcup_{i\in I} X_i \subseteq Y\}\end{aligned}$$

for any non-empty family $\{X_i\}_{i\in I} \subseteq \mathcal{X}$, and in fact a proof of this claim is supplied in the form of Corollary 2.32 in [13]. From the second equality, we see that the join operator does not coincide with union.

Conveniently, the map that takes an arbitrary observation about D to its closest approximation in $\mathcal{D}$ is the closure operator $C_{\mathcal{D}}$ induced on D by $\mathcal{D}$. Definition 3.3 and Proposition 3.4 are derived from specialisations of Definition 7.1 and Proposition 7.2 in [13] to the power set $\mathcal{P}(X)$ of a set X. The proof of Proposition 3.4 is straightforward, and we omit it here.

Definition 3.3. Let X be a set. A map $C : \mathcal{P}(X) \longrightarrow \mathcal{P}(X)$ is called a *closure operator* (on X) if, for any $U, V \subseteq X$,

CL1: $U \subseteq C(U)$

CL2: $U \subseteq V \Rightarrow C(U) \subseteq C(V)$

CL3: $C(C(U)) = C(U)$

A set $U \subseteq X$ is called *closed* if $U = C(U)$. Given a closure operator C, the set of all closed subsets of X under C is simply $\{U \subseteq X \mid U = C(U)\}$.

Proposition 3.4. *Let X be a set and let C be a closure operator on $(\mathcal{P}(X), \subseteq)$.*

i) The set of all closed subsets of X is $\mathcal{X} = \{C(U) \mid U \subseteq X\}$, and $X \in \mathcal{X}$.

ii) For any $U \subseteq X$, $C(U) = \bigcap\{V \in \mathcal{X} \mid U \subseteq V\}$

Topped $\bigcap$-structures and closure operators are closely connected, as shown by the following theorem, derived from Theorem 7.3 in [13]. As before, the proof of this result is straightforward, and we omit it here.

Theorem 3.5. *Let X be a set, and let $\mathcal{X}$ be a topped $\bigcap$-intersection structure on X. The map*

$$C_{\mathcal{X}} : \mathcal{P}(X) \longrightarrow \mathcal{P}(X) : U \mapsto \bigcap\{V \in \mathcal{X} \mid U \subseteq V\}$$

is a closure operator on X. Let $C : \mathcal{P}(X) \longrightarrow \mathcal{P}(X)$ be a closure operator on X. Then the family

$$\mathcal{X}_C = \{U \subseteq X \mid U = C(U)\}$$

of closed subsets of X under C is a topped $\bigcap$-structure on X. Furthermore, $(\mathcal{X}_C, \subseteq)$ is a complete lattice in which for any family $\{X_i\}_{i \in I} \subseteq \mathcal{X}_C$,

$$\begin{aligned} \bigwedge_{i \in I} X_i &= \bigcap_{i \in I} X_i \\ \bigvee_{i \in I} X_i &= C(\bigcup_{i \in I} X_i) \end{aligned}$$

We shall invoke the duality between sets and properties and treat the observations in the topped $\bigcap$-structure $\mathcal{D}$ as properties of the members of D. That is, we treat each $U \in \mathcal{D}$ as a property θ_U, and if θ_U is true of $d \in D$, we write $d \models \theta_U$, so that $U = \{d \in D \mid d \models \theta_U\}$. This gives rise to a family $\Theta_{\mathcal{D}} = \{\theta_U\}_{U \in \mathcal{D}}$ of properties. We let

$$D_d = \{U \in \mathcal{D} \mid d \in U\}$$

denote the family of observations about a given $d \in D$. As expected, the corresponding set of properties is just

$$\Theta_d = \{\theta \in \Theta_{\mathcal{D}} \mid d \models \theta\} = \{\theta_U \in \Theta_{\mathcal{D}} \mid U \in D_d\}$$

We shall sometimes exploit the duality between sets and properties directly and write a property θ as a subset of D.

To order the members of D, for any $c, d \in D$ we shall take d to be stronger or more refined than c if every observation that is true of c is also true of d.

Definition 3.6. For each $d \in D$, let $D_d = \{U \in \mathcal{D} \mid d \in U\}$, and let $h : D \longrightarrow \mathcal{P}(\mathcal{D}) : d \mapsto D_d$. We define the binary relation $\preceq$ on D such that for any $c, d \in D$,

$$c \preceq d \quad \text{if and only if} \quad h(c) \subseteq h(d)$$

It follows readily from Definition 3.6 that $\preceq$ is a quasi-order on D. We can recover a partial order on D, which we will also denote by $\preceq$, by taking $c \approx d$ if and only if $c \preceq d$ and $d \preceq c$. From Definition 3.1(ii) however, for no two distinct elements $c, d \in D$ is it the case that $c \approx d$, because then $h(c) = h(d)$, which violates Definition 3.1(ii).

The following result is an easy consequence of Definition 3.6.

Proposition 3.7. *Let $\mathcal{D}_0$ be an admissible set of observations about a set D, with $d_\mathbf{0}, d_\mathbf{1} \in D$, and let $\mathcal{D}$ be the corresponding topped $\bigcap$-structure. Let $h : D \longrightarrow \mathcal{P}(\mathcal{D})$ be such that $h(d) = \{U \in \mathcal{D} \mid d \in U\}$. Then $\bigcap_{d \in D} h(d) = \{D\}$ and $\bigcup_{d \in D} h(d) = \mathcal{D}$.*

Proof. From Definition 3.6, for any $U \in \bigcap_{d \in D} h(d)$ we have

$$U \in \bigcap_{d \in D} h(d) \;\Leftrightarrow\; \forall d \in D.[U \in h(d)] \;\Leftrightarrow\; \forall d \in D.[d \in U] \;\Leftrightarrow\; D \subseteq U$$

and since $D = \bigcup \mathcal{D}$, $U \subseteq D$ so that $U = D$ and hence $\bigcap_{d \in D} h(d) = \{D\}$. For the second part, we have already that $\bigcup_{d \in D} h(d) \subseteq \mathcal{D}$. Since $\mathcal{D}_0$ is an admissible set of observations about D, it follows that $D = \bigcup \mathcal{D}$. Hence for any $U \in \mathcal{D}$ there is $d \in D$ with $d \in U$, so that $U \in h(d)$. It follows that $\mathcal{D} \subseteq \bigcup_{d \in D} h(d)$, and hence that $\bigcup_{d \in D} h(d) = \mathcal{D}$. □

Corollary 3.8. *Let $\mathcal{D}_0$ be an admissible set of observations about a set D, with $d_{\mathbf{0}}, d_{\mathbf{1}} \in D$, and let $\mathcal{D}$ be the corresponding topped $\bigcap$-structure on D. Let $D_d = \{U \in \mathcal{D} \mid d \in U\}$, let $h : D \longrightarrow \mathcal{P}(\mathcal{D}) : d \mapsto D_d$ and let $\preceq \subseteq D \times D$ be such that $c \preceq d$ if and only if $h(c) \subseteq h(d)$. Then, under $\preceq$, $d_{\mathbf{1}}$ and $d_{\mathbf{0}}$ are, respectively, the top and bottom elements of D.*

Proof. For any $d \in D$, $h(d) \subseteq \mathcal{D}$, and since $h(d_{\mathbf{1}}) = \mathcal{D}$ it follows that $d \preceq d_{\mathbf{1}}$. Suppose that there is $d \in D$ such that $d \neq d_{\mathbf{1}}$ and $d_{\mathbf{1}} \preceq d$. Then $h(d_{\mathbf{1}}) \subseteq h(d)$, and it follows that $h(d) = \mathcal{D}$. For every $U \in \mathcal{D}_0$ we then have $d, d_{\mathbf{1}} \in U$, in which case $\mathcal{D}_0$ is not an admissible set of observations about D and we reach a contradiction. Hence $d_{\mathbf{1}}$ is the top element of D. By Proposition 3.7, for any $d \in D$ we have $\{D\} \subseteq h(d)$, and since $h(d_{\mathbf{0}}) = \{D\}$ it follows that $d_{\mathbf{0}} \preceq d$. Suppose that there is $d \in D$ with $d \neq d_{\mathbf{0}}$ and $d \preceq d_{\mathbf{0}}$. Then $h(d) \subseteq \{D\}$, and hence $h(d) = \{D\}$. It follows that $\mathcal{D}_0$ is not an admissible set of observations, and we again reach a contradiction. Hence $d_{\mathbf{0}}$ is the bottom element of D. □

For $c, d \in D$, if $c \preceq d$ we may think of d as representing a stronger property than c.

Proposition 3.9. *For any $c, d \in D$, if $c \preceq d$ then $C_{\mathcal{D}}(\{d\}) \subseteq C_{\mathcal{D}}(\{c\})$.*

Proof. For any $c, d \in D$,

$$\begin{array}{rll} c \preceq d \Rightarrow & h(c) \subseteq h(d) & \text{(definition of } \preceq\text{)} \\ \Rightarrow & \{U \in \mathcal{D} \mid c \in U\} \subseteq \{U \in \mathcal{D} \mid d \in U\} & \text{(definition of } h\text{)} \\ \Rightarrow & \bigcap\{U \in \mathcal{D} \mid d \in U\} \subseteq \bigcap\{U \in \mathcal{D} \mid c \in U\} & \text{(Lemma 2.22(v) in [13])} \\ \Rightarrow & C_{\mathcal{D}}(\{d\}) \subseteq C_{\mathcal{D}}(\{c\}) & \text{(definition of } C_{\mathcal{D}}\text{)} \end{array}$$

which gives us the result, as required. □

The members of $\Theta_{\mathcal{D}}$ may be ordered in a way that reflects the order induced on D by $\mathcal{D}$. To do so, we lift the existing order on D to a power order. From Definition 2.30 in [8] we have the following.

Definition 3.10. For any set X and any binary relation $R \subseteq X \times X$, we define for all $U, V \subseteq X$ the *lower power relation* R^0, the *upper power relation* R^1 and the *full power relation* R^+, where $R^0, R^1, R^+ \subseteq \mathcal{P}(X) \times \mathcal{P}(X)$, by

$$\begin{array}{rll} UR^0V & \text{if and only if} & \forall u \in U.\exists v \in V.[uRv] \\ UR^1V & \text{if and only if} & \forall v \in V.\exists u \in U.[uRv] \\ UR^+V & \text{if and only if} & UR^0V \text{ and } UR^1V \end{array}$$

The lower and upper power orders $\preceq^0$ and $\preceq^1$ are not well-behaved with respect to the top and bottom elements of D. To illustrate, suppose that we have $d_{\mathbf{1}} \in U \cap V$ for $U, V \subseteq D$. Then $U \preceq^0 V$ and $V \preceq^0 U$. Similarly, if $U, V \subseteq D$ with $d_{\mathbf{0}} \in U \cap V$ then $U \preceq^1 V$ and $V \preceq^1 U$. To prevent subsets of D from being trivially ordered in this way, we use the full power order $\preceq^+$ and for any $\phi, \psi \in \Theta_{\mathcal{D}}$ we therefore define

$$\phi \sqsubseteq \psi \quad \text{if and only if} \quad \{d \in D \mid d \models \phi\} \preceq^+ \{d \in D \mid d \models \psi\}$$

To $\Theta_{\mathcal{D}}$ we add the property $\theta_{\mathbf{0}} = \{d_{\mathbf{0}}\}$, and we write $\theta_{\mathbf{1}}$ for $\{d_{\mathbf{1}}\}$.

Proposition 3.11. *The elements $\theta_{\mathbf{1}}$ and $\theta_{\mathbf{0}}$ are, respectively the top and bottom elements of $\Theta_{\mathcal{D}}$.*

Proof. For any $\theta \in \Theta_{\mathcal{D}}$, let $U_\theta = \{d \in D \mid d \models \theta\}$. Trivially, we have

$$\forall c \in U_\theta.\exists d \in \{d_\mathbf{1}\}.[c \preceq d] \text{ and } \forall d \in \{d_\mathbf{1}\}.\exists c \in U_\theta.[c \preceq d]$$

so that $\theta \sqsubseteq \theta_\mathbf{1}$. Suppose that there is $\theta \in \Theta_{\mathcal{D}}$ with $\theta \neq \theta_\mathbf{1}$ and $\theta_\mathbf{1} \sqsubseteq \theta$. Then $\{d_\mathbf{1}\} \preceq^+ U_\theta$, in which case there is $d \in D$ with $d \neq d_\mathbf{1}$ and $d_\mathbf{1} \preceq d$, and we reach a contradiction. It follows that $\theta_\mathbf{1}$ is the top element of $\Theta_{\mathcal{D}}$. Similarly, we can show that $\theta_\mathbf{0}$ is the bottom element of $\Theta_{\mathcal{D}}$, which gives us the result as required. □

A power order is a quasi-order even if the order from which it is derived is a partial order. For example, if $(P, \leq)$ is a partially ordered set and $x, y, z \in P$ with $x \leq y \leq z$, then $\{x, z\} \leq^0 \{y, z\}$ and $\{y, z\} \leq^0 \{x, z\}$ even though $\{x, z\} \neq \{y, z\}$. We recover a partial order, which we also denote by $\sqsubseteq$, by taking $\theta \approx \gamma$ if and only if $\theta \sqsubseteq \gamma$ and $\gamma \sqsubseteq \theta$.

We shall refer to a subset $\Theta \subseteq \Theta_{\mathcal{D}}$ as a specification. From the definition of $\sqsubseteq$, for any $\phi, \psi \in \Theta_{\mathcal{D}}$, if $\phi \sqsubseteq \psi$ then, based on Proposition 3.9, we may think of ψ as a stronger property than ϕ. The specification $\uparrow\!\phi$ may then be thought of as comprising those observations that are at least as strong as ϕ. Correspondingly, $(\uparrow\phi)^c = \Theta_{\mathcal{D}} \setminus \uparrow\phi$ may be thought of as those observations, none of which is stronger than ϕ. Consequently, let $\mathcal{S}_0 = \{\uparrow\phi \mid \phi \in \Theta_{\mathcal{D}}\} \cup \{(\uparrow\phi)^c \mid \phi \in \Theta_{\mathcal{D}}\}$ be a non-empty family of subsets of $\Theta_{\mathcal{D}}$. With the development of Section 3.2 in mind, we close $\mathcal{S}_0$ under non-empty intersections and arbitrary unions to form the family $\mathcal{S}_{\mathcal{D}}$ of subsets of $\Theta_{\mathcal{D}}$. Observe that $\mathcal{S}_{\mathcal{D}}$ is a topped $\bigcap$-structure $\mathcal{S}_{\mathcal{D}}$ on $\Theta_{\mathcal{D}}$, and that $\emptyset \in \mathcal{S}_{\mathcal{D}}$. We denote by C_S the closure operator induced on $\Theta_{\mathcal{D}}$ by $\mathcal{S}_{\mathcal{D}}$.

Remark 3.12. The family $\mathcal{S}_{\mathcal{D}}$ was developed with the logic of Section 3.2 in mind. In particular, arbitrary unions were needed so that the lattice of surrogate degrees of disposition that we will develop in Section 3.1.2 would be (finitely) distributive. In turn, this property was required so that condition **Dn2** of Definition 3.49 would be satisfied by our proposed disjunction (see especially Proposition 3.53). Intuitively, if the negation, conjunction and disjunction connectives of the logic developed in Section 3.2 were required to satisfy different axioms, or if one or other connective were omitted, $\mathcal{S}_{\mathcal{D}}$ could be formed differently. For example, if only conjunction were required, we could choose $\mathcal{S}_{\mathcal{D}}$ to be the family of closed subsets of $\Theta_{\mathcal{D}}$ that arises from the closure operator $C_S : \mathcal{P}(\Theta_{\mathcal{D}}) \longrightarrow \mathcal{P}(\Theta_{\mathcal{D}}) : \Theta \mapsto \Theta^{lu}$.

We shall treat specifications conjunctively, so $d \in D$ satisfies $\Theta \subseteq \Theta_{\mathcal{D}}$ if for every $\theta \in \Theta$, $d \models \theta$. The specification Θ thus gives rise to the set

$$D_\Theta = \bigcap_{\theta \in \Theta} \{d \in D \mid d \models \theta\} = \{d \in D \mid \Theta \subseteq \Theta_d\}$$

of degrees of disposition.

3.1.2 Surrogate Degrees of Disposition

From Proposition 3.4(ii), given an arbitrary specification $\Theta \subseteq \Theta_{\mathcal{D}}$, the smallest closed specification at our disposal is simply

$$P_\Theta \;=\; \bigcap\{U \in \mathcal{S}_{\mathcal{D}} \mid \Theta \subseteq U\} \;=\; C_S(\Theta)$$

In contrast, a closed specification $U \in \mathcal{S}_{\mathcal{D}}$ such that $U \subseteq \Theta^c$, where $\Theta^c = \Theta_{\mathcal{D}} \setminus \Theta$, contains properties that are excluded from Θ, and so may be thought of as being converse to Θ. The largest such closed specification that we can construct is then just

$$Q_\Theta \;=\; \bigcup\{U \in \mathcal{S}_{\mathcal{D}} \mid U \subseteq \Theta^c\}$$

The definition of P_Θ and Q_Θ is an adaption of an idea presented in [9] (compare the definition of the p- and q-sets on p338 in that work). We may think of P_Θ as a closed specification for the degree of disposition held by an agent towards a given formula. Correspondingly, we may think of Q_Θ as a closed specification for the degree of *counter*-disposition held by an agent towards the same formula.

In general, we expect that $D_{P_\Theta} \cap D_{Q_\Theta} \neq \emptyset$, so it is possible for P_Θ and Q_Θ to embody an element of ambiguity. We shall use this ambiguity to characterise the surrogate degrees of disposition that we derive from $\mathcal{S}_\mathcal{D}$. To derive these surrogates, we adapt an idea presented in [15] (compare Example 7.2 on p229 in [12] for an alternative formulation of what we present here).

Definition 3.13. A *$\mathcal{D}$-surrogate* is a pair (U, V) where $U, V \in \mathcal{S}_\mathcal{D}$. It is

i) *overdefined* if $U \cap V \neq \emptyset$

ii) *consistent* if $U \cap V = \emptyset$

iii) *exact* if it is consistent and furthermore $U \cup V = \Theta_\mathcal{D}$

The set of all $\mathcal{D}$-surrogates is denoted by $\mathfrak{T}_\mathcal{D}$. For an arbitrary specification $\Theta \subseteq \Theta_\mathcal{D}$, we define P_Θ and Q_Θ to be the subsets of $\Theta_\mathcal{D}$ given by

$$\begin{aligned} P_\Theta &= \bigcap\{U \in \mathcal{S}_\mathcal{D} \mid \Theta \subseteq U\} \\ Q_\Theta &= \bigcup\{U \in \mathcal{S}_\mathcal{D} \mid U \subseteq \Theta^c\} \end{aligned}$$

and the corresponding $\mathcal{D}$-surrogate is (P_Θ, Q_Θ). For $\Theta, \Upsilon \subseteq \Theta_\mathcal{D}$, the $\mathcal{D}$-surrogate of the pair (Θ, Υ) is given by (P_Θ, Q_Υ).

Given a $\mathcal{D}$-surrogate $t \in \mathfrak{T}_\mathcal{D}$, we will write $t = (P_t, Q_t)$ when the specifications that gave rise to P_t and Q_t are not supplied and to emphasise the role of the P- and Q-sets in describing degrees of disposition and counter-disposition towards a given sentence $x \in X$.

Because P_t and Q_t are sets, $\mathcal{D}$-surrogates can be ordered by inclusion $\subseteq$. In [15], two possible orders are suggested (compare also p11 in [11]):

i) the *knowledge* order $\leq_k$, under which $s \leq_k t$ if $P_s \subseteq P_t$ and $Q_s \subseteq Q_t$

ii) the *truth* order $\leq_t$, under which $s \leq_t t$ if $P_s \subseteq P_t$ and $Q_t \subseteq Q_s$

where s and t are $\mathcal{D}$-surrogates. For our work, we will use the truth order, which we denote simply by $\leq_T$. Trivially, $\leq_T$ is a partial order (this property is inherited via the componentwise application of the inclusion order). Under the truth order, there are four $\mathcal{D}$-surrogates in $(\mathfrak{T}_\mathcal{D}, \leq_T)$ that are of particular significance to us, *viz.*

i) $(\emptyset, \emptyset)$, which indicates that the agent holds neither disposition nor counter-disposition towards a sentence $x \in X$. We shall view this value as representing an "indeterminate" degree of disposition, and we will use it to indicate that the agent holds an "unknown" degree of disposition towards x.

ii) $(\emptyset, \Theta_\mathcal{D})$, which indicates that the agent does not hold the disposition at all towards x, and is the least element $\mathbf{0}$ in $\mathfrak{T}_\mathcal{D}$.

iii) $(\Theta_\mathcal{D}, \emptyset)$, which indicates that the agent holds the disposition maximally towards x. This element is the greatest element $\mathbf{1}$ in $\mathfrak{T}_\mathcal{D}$.

iv) $(\Theta_{\mathcal{D}}, \Theta_{\mathcal{D}})$, which indicates that the agent holds a fully contradictory disposition towards x.

Definition 3.14. Let $\mathsf{D}(\mathfrak{T}_{\mathcal{D}})$ be the set of $\mathcal{D}$-surrogates in $\mathfrak{T}_{\mathcal{D}}$ defined by

$$\mathsf{D}(\mathfrak{T}_{\mathcal{D}}) \;=\; \{t \in \mathfrak{T}_{\mathcal{D}} \mid P_t = Q_t\}$$

We call the set $\mathsf{D}(\mathfrak{T}_{\mathcal{D}})$ the *diagonal* of $\mathfrak{T}_{\mathcal{D}}$.

A $\mathcal{D}$-surrogate $t \in \mathsf{D}(\mathfrak{T}_{\mathcal{D}})$, represents a state in which the agent is unsure of its degree of disposition towards a sentence $x \in X$. These $\mathcal{D}$-surrogates do not convey any information about the degree of disposition of the agent towards x because we cannot tell whether the agent is more or less disposed than counter-disposed towards x, and so we may view them as being equivalent to the "unknown" degree of disposition represented by the $\mathcal{D}$-surrogate $(\emptyset, \emptyset)$. The values of the diagonal range from $(\emptyset, \emptyset)$ through increasing degrees of contradiction until we reach $(\Theta_{\mathcal{D}}, \Theta_{\mathcal{D}})$.

We define meet and join so as to respect the topped $\bigcap$-structure $(\mathcal{S}_{\mathcal{D}}, \subseteq)$ on $\Theta_{\mathcal{D}}$.

Definition 3.15. For any $s, t \in \mathfrak{T}_{\mathcal{D}}$, we define meet $\wedge$ and join $\vee$ such that

i) $(P_s, Q_s) \wedge (P_t, Q_t) = (P_s \cap P_t, Q_s \cup Q_t)$

ii) $(P_s, Q_s) \vee (P_t, Q_t) = (P_s \cup P_t, Q_s \cap Q_t)$

For a non-empty subset T of $\mathfrak{T}_{\mathcal{D}}$, we then have

$$\begin{aligned} \bigwedge T &= \inf T = (\bigcap_{t\in T} P_t, \bigcup_{t\in T} Q_t) \\ \bigvee T &= \sup T = (\bigcup_{t\in T} P_t, \bigcap_{t\in T} Q_t,) \end{aligned}$$

where for each $t \in T$, (P_t, Q_t) is the corresponding $\mathcal{D}$-surrogate. Since $\bigcap_{t\in T} P_t$ and $\bigcup_{t\in T} Q_t$ are both closed sets in $\Theta_{\mathcal{D}}$, it follows that $\bigwedge T \in \mathfrak{T}_{\mathcal{D}}$, and similarly for $\bigvee T$. Furthermore, because $\mathfrak{T}_{\mathcal{D}}$ contains a top element $(\Theta_{\mathcal{D}}, \emptyset)$, by Theorem 2.31 in [13] $\mathfrak{T}_{\mathcal{D}}$ is a complete lattice. We state this as the following result.

Proposition 3.16. *The partially ordered set* $(\mathfrak{T}_{\mathcal{D}}, \leq_T)$ *is a complete lattice.*

As an easy consequence, we then have the following (compare Definition 4.4 in [13]).

Proposition 3.17. *The lattice* $(\mathfrak{T}_{\mathcal{D}}, \leq_T)$ *is distributive.*

Proof. Applying Definition 3.15, for any $r, s, t \in \mathfrak{T}_{\mathcal{D}}$,

$$\begin{aligned} r \wedge (s \vee t) &= (P_r, Q_r) \wedge ((P_s, Q_s) \vee (P_t, Q_t)) \\ &= (P_r, Q_r) \wedge (P_s \cup P_t, Q_s \cap Q_t) \\ &= (P_r \cap (P_s \cup P_t), Q_r \cup (Q_s \cap Q_t)) \\ &= ((P_r \cap P_s) \cup (P_r \cap P_t), (Q_r \cup Q_s) \cap (Q_r \cup Q_t)) && \text{(set theory)} \\ &= (P_r \cap P_s, Q_r \cup Q_s) \vee (P_r \cap P_t, Q_r \cup Q_t) \\ &= ((P_r, Q_r) \wedge (P_s, Q_s)) \vee ((P_r, Q_r) \wedge (P_t, Q_t)) \\ &= (r \wedge s) \vee (r \wedge t) \end{aligned}$$

which gives us the result as required. □

In a lattice-theoretic setting, different types of complementation are commonly used to derive a negation operation. For example, in a lattice with top and bottom elements $\mathbf{0}$ and $\mathbf{1}$, an element y is called a complement of an element x if $x \wedge y = \mathbf{0}$ and $x \vee y = \mathbf{1}$ (compare Definition 4.13 in [13]). Given $s = (P_s, Q_s) \in \mathfrak{T}_{\mathcal{D}}$, from Definition 3.15 a complement $t = (P_t, Q_t)$ of s would need to have the form $(\Theta_D \setminus P_s, \Theta_D \setminus Q_s)$ to satisfy $s \wedge t = \mathbf{0}$ and $s \vee t = \mathbf{1}$. These set-theoretic complements are, however, not guaranteed to be closed sets of Θ_D, so we turn instead to a weaker notion of complement called involution (compare Definition 5.18 in [12]).

Definition 3.18. Let $(X, \leq)$ be an ordered set. An order-reversing map $' : X \to X$ is called an *involution* if for any $x, y \in X$,

i) $x \leq y$ implies $y' \leq x'$

ii) $x'' = x$

For any $x \in X$, we shall overload terminology and also refer to x' as the *involution of* x.

We shall effect an involution on $\mathfrak{T}_{\mathcal{D}}$ by reversing the roles of the P and Q sets within a given $\mathcal{D}$-surrogate. That is, given $t = (P_t, Q_t)$, we define a unary operation $'$ on $\mathfrak{T}_{\mathcal{D}}$ such that $t' = (Q_t, P_t)$.

Proposition 3.19. *The operation $'$ is an involution.*

Proof. For any two $\mathcal{D}$-surrogates $s = (P_s, Q_s)$ and $t = (P_t, Q_t)$ in $\mathfrak{T}_{\mathcal{D}}$, if $s \leq_T t$ then $P_s \subseteq P_t$ and $Q_t \subseteq Q_s$, and we then have $t' = (Q_t, P_t) \leq_T (Q_s, P_s) = s'$. Furthermore, $s' = (Q_s, P_s)$ and hence $s'' = (Q_s, P_s)' = (P_s, Q_s) = s$. □

Note that for $t \in \mathsf{D}(\mathfrak{T}_{\mathcal{D}})$, $t' = t$. Furthermore, for $s, t \in \mathfrak{T}_{\mathcal{D}}$, if s is such that $P_s \subseteq P_t \cap Q_t$ and $P_t \cup Q_t \subseteq Q_s$, then $s \leq_T t, t'$. We shall need to account for this possibility in the derivation of a logic for X in Section 3.2 (see Definition 3.27 in particular).

Henceforth, for $\mathfrak{T}_{\mathcal{D}}$ we shall denote meet by $*$ and join by $+$ (and correspondingly, we denote $\bigwedge$ by $\prod$ and $\bigvee$ by $\sum$), so that $(\mathfrak{T}_{\mathcal{D}}, +, *, ')$ is an algebra of type $\mathcal{F}$. With the partial order $\leq_T$, we then have a bounded r-algebra $(\mathfrak{T}_{\mathcal{D}}, +, *, ', \mathbf{0}, \mathbf{1}, \leq_T)$ of type $\mathcal{F}$.

Definition 3.20. Let D be a non-empty set, let $\mathcal{D}_0$ be an admissible set of observations about D and let $d_\mathbf{0}$ and $d_\mathbf{1}$ exist in D. A *disposition* is a triple $\mathbf{D} = (D, \mathcal{D}_0, \mathbf{T})$, where $\mathbf{T} = (\mathfrak{T}_{\mathcal{D}}, +, *, ', \mathbf{0}, \mathbf{1}, \leq_T)$ is a bounded r-algebra such that

i) $\mathcal{D}$ is the topped $\bigcap$-structure on D formed by closing $\mathcal{D}_0$ under non-empty intersections

ii) $\mathfrak{T}_{\mathcal{D}}$ is the corresponding set of $\mathcal{D}$-surrogates

iii) $\leq_T$ is the truth-order, $*$ and $+$ represent the meet and join operations with respect to $\leq_T$, and $'$ is a unary operation on $\mathfrak{T}_{\mathcal{D}}$ such that for any $t = (P_t, Q_t) \in \mathfrak{T}_{\mathcal{D}}$, $t' = (Q_t, P_t)$.

3.2 A Logic of Disposition

Let us presuppose a disposition $\mathbf{D} = (D, \mathcal{D}_0, \mathbf{T})$, where $\mathbf{T} = (\mathfrak{T}_{\mathcal{D}}, +, *, ', \mathbf{0}, \mathbf{1}, \leq_T)$ as in Definition 3.20. Given the sentence algebra (X, F) of type $\mathcal{F}$ over X_0, we may use the reference assignment v to derive the order $\leq_X$ and Alexandroff topology ΩX on X from $\mathbf{T}$ as described in Section 2.2.2. By including a sheaf H of $\mathbf{T}$-simulations, we then form the ordered r-algebraic space $(X, H, \leq_X)$ of type $\mathcal{F}$ over $\mathbf{T}$.

We shall use $\mathbf{T}$ to derive a logic for (X, F). From [12], the principle constituent of a system of logic is the relation of consequence, which also distinguishes a logic from an algebra. We therefore commence our development with the notion of consequence. We shall use [12] as a guide, and many of the definitions and results presented here may be found there, although we have adapted the notation used to match our conventions.

3.2.1 Logical Consequence

From [12], a consequence function for a logic on a set A of sentences is a function $C_n : \mathcal{P}(A) \to \mathcal{P}(A)$ that associates a set U of sentences with a set $C_n(U)$, the logical consequences of U, and induces an equivalence relation $\simeq_{C_n}$ on A such that for any $x, y \in A$ we have

$$x \simeq_{C_n} y \quad \text{if and only if} \quad C_n(\{x\}) = C_n(\{y\})$$

which extends to arbitrary subsets S and T of A as

$$S \simeq_{C_n} T \quad \text{if and only if} \quad C_n(S) = C_n(T)$$

When we combine C_n with an algebra of a specified type, we obtain a *prelogic*. If C_n satisfies certain additional conditions, we obtain a logic. Definition 3.21 is adapted from Definitions 3.2(ii), 3.3(i) and 3.7 in [12].

Definition 3.21. Let $\mathcal{F}$ be a language of algebras.

1. A *prelogic* of type $\mathcal{F}$ is a pair $(\mathbf{A}, C_n)$ where $\mathbf{A}$ is an algebra of type $\mathcal{F}$ and $C_n : \mathcal{P}(A) \to \mathcal{P}(A)$ is a set-to-set function. We denote the prelogic by $\mathfrak{L}_A^-$ and call $\mathbf{A}$ the *algebra of sentences* of $\mathfrak{L}_A^-$, C_n the *consequence function* of $\mathfrak{L}_A^-$ and $\simeq_{C_n}$ the *logical equivalence* on $\mathfrak{L}_A^-$, also denoted by $\leftrightarrow$. If C_n satisfies the conditions

   L1: For any $U \subseteq A$, $U \subseteq C_n(U)$

   L2: For any $U \subseteq A$, $C_n(C_n(U)) \subseteq C_n(U)$

   L3: For any $U, V \subseteq A$, $U \subseteq V \Rightarrow C_n(U) \subseteq C_n(V)$

   then $(\mathbf{A}, C_n)$ is a logic of type $\mathcal{F}$, denoted by $\mathfrak{L}_A$. For $x \in A$, x is a *logical consequence* of $U \subseteq A$ if $x \in C_n(U)$, and we write $U \vdash_{\mathfrak{L}_A} x$. The set of theorems of $\mathfrak{L}_A$ is given by $Th(\mathfrak{L}_A) = C_n(\emptyset)$.

2. Let $\mathfrak{L}_A = (\mathbf{A}, C_A)$ and $\mathfrak{L}_B = (\mathbf{B}, C_B)$ be two logics of type $\mathcal{F}$, and let $f : A \longrightarrow B$ be a homomorphism from $\mathbf{A}$ to $\mathbf{B}$.

   i) If f maps A into (onto) B, then f is a *weak homomorphism of* $\mathfrak{L}_A$ *into (onto)* $\mathfrak{L}_B$ if for any $U \subseteq A$, $C_A(U) \subseteq f^{\leftarrow}(C_B(f(U)))$, and a *strong homomorphism of* $\mathfrak{L}_A$ *into (onto)* $\mathfrak{L}_B$ if $C_A(U) = f^{\leftarrow}(C_B(f(U)))$.

   ii) If f is an epimorphism from $\mathbf{A}$ to $\mathbf{B}$, then $\mathfrak{L}_B$ is a *weak (strong) homomorphic image* of $\mathfrak{L}_A$ if f is a weak (strong) homomorphism of $\mathfrak{L}_A$ onto $\mathfrak{L}_B$.

   iii) If f is an isomorphism from $\mathbf{A}$ to $\mathbf{B}$, then f is an isomorphism of $\mathfrak{L}_A$ to $\mathfrak{L}_B$ if it is a strong homomorphism of $\mathfrak{L}_A$ onto $\mathfrak{L}_B$.

From Conditions **L1** and **L3**, $C_n(U) \subseteq C_n(C_n(U))$, and in conjunction with Condition **L2** we have $C_n(U) = C_n(C_n(U))$. Thus we recognise in Definition 3.21 the conditions of Definition 3.3, which allows us to use a closure operator on X as a consequence function.

Let C_n be a consequence function for the algebra (X, F) of sentences, so that $\mathfrak{L}_X = (\mathbf{X}, C_n)$ is a logic of type $\mathcal{F}$. We would like C_n to preserve the degree of disposition in the sense that for any $U \subseteq X$ and $x \in X$,

$$U \vdash_{\mathfrak{L}_X} x \quad \text{if and only if} \quad \forall u \in [X_0 \longrightarrow \mathfrak{T}_{\mathcal{D}}].\forall t \in \mathfrak{T}_{\mathcal{D}}.[\forall y \in U.[t \leq_T u(y)] \Rightarrow t \leq_T u(x)]$$

(compare the identity (5) in [16]). One means of accomplishing this preservation is to fix a topped $\bigcap$-structure $\mathcal{C}_T$ on $\mathfrak{T}_{\mathcal{D}}$ and then define $\vdash_{\mathfrak{L}_X}$ to be such that

$$U \vdash_{\mathfrak{L}_X} x \quad \text{if and only if} \quad \forall u \in [X_0 \longrightarrow \mathfrak{T}_{\mathcal{D}}].[u(x) \in C_T(u(U))]$$

where C_T is the closure operator on $\mathfrak{T}_{\mathcal{D}}$ induced by $\mathcal{C}_T$ (compare p204 in [12]; see also the statement (5.2) in [17]). Intuitively, for any $U \subseteq X$ we could then set $C_n(U) = \{x \in X \mid U \vdash_{\mathfrak{L}_X} x\}$.

For work we have yet to complete, we shall require that C_n should be an algebraic closure operator. Recall the notion of a directed set from Definition 2.13. Where this is applied to a non-empty family of sets, we have a directed family of sets in which the directed joins are referred to as directed unions. Definitions 3.22, 3.23 and Theorem 3.24 are respectively based on Definitions 7.10, 7.12 and Theorem 7.14 in [13]. The proof of Theorem 3.24 is supplied in [13] and so we omit it here.

Definition 3.22. Let $\mathcal{A}$ be a non-empty family of subsets of a non-empty set A. The family $\mathcal{A}$ is *closed under directed unions* if $\bigcup_{i \in I} A_i \in \mathcal{A}$ for any directed family $\{A_i\}_{i \in I}$ in $\mathcal{A}$. It is an *algebraic* $\bigcap$*-structure* if

i) $\bigcap_{i \in I} A_i \in \mathcal{A}$ for any non-empty family $\{A_i\}_{i \in I}$ in $\mathcal{A}$

ii) $\bigcup_{i \in I} A_i \in \mathcal{A}$ for any directed family $\{A_i\}_{i \in I}$ in $\mathcal{A}$

An algebraic $\bigcap$-structure $\mathcal{A}$ on a set A is thus an $\bigcap$-structure on A that is closed under directed unions. Intuitively, a topped algebraic $\bigcap$-structure on A is an algebraic $\bigcap$-structure $\mathcal{A}$ on A with $A \in \mathcal{A}$.

Definition 3.23. A closure operator C on a set A is *algebraic* if, for all $U \subseteq A$,

$$C(U) \quad = \quad \bigcup \{C(V) \mid V \subseteq U \text{ and } V \text{ is finite}\}$$

Theorem 3.24. *Let C be a closure operator on a set A, and let $\mathcal{A}_C$ be the associated topped $\bigcap$-structure. Then the following are equivalent:*

i) C is an algebraic closure operator

ii) for every directed family $\{A_i\}_{i \in I}$ of subsets of A,

$$C(\bigcup_{i \in I} A_i) \quad = \quad \bigcup_{i \in I} C(A_i)$$

iii) $\mathcal{A}_C$ is an algebraic $\bigcap$-structure

From Definition 2.20 in [13], we have the following.

Definition 3.25. Let $(A, \leq)$ be a lattice. A non-empty subset B of A is called a *filter (of A)* if

i) $x, y \in B$ implies $x \wedge y \in B$

ii) $x \in B$, $y \in A$ and $x \leq y$ imply $y \in B$

The set of all filters of A is denoted by $\mathfrak{F}(A)$ and carries the usual inclusion order.

Proposition 3.26. *Let $(A, \leq)$ be a lattice. Then $\mathfrak{F}(A)$ is a topped algebraic $\bigcap$-structure on A.*

Proof. Let $\{B_i\}_{i\in I}$ be a non-empty family of filters of A, and let $B = \bigcap_{i\in I} B_i$. Applying Definition 3.25, we then have

$$x, y \in B \;\Rightarrow\; \forall i \in I.[x, y \in B_i] \;\Rightarrow\; \forall i \in I.[x \wedge y \in B_i] \;\Rightarrow\; x \wedge y \in \bigcap_{i\in I} B_i \;\Rightarrow\; x \wedge y \in B$$

Next, suppose that $x \in B$, $y \in A$ and $x \leq y$. Applying Definition 3.25 once more, we then have

$$x \in B \;\Rightarrow\; \forall i \in I.[x \in B_i] \;\Rightarrow\; \forall i \in I.[y \in B_i] \;\Rightarrow\; y \in \bigcap_{i\in I} B_i \;\Rightarrow\; y \in B$$

so that B is also a filter. Now let $\mathcal{B} = \{B_i\}_{i\in I}$ be a directed family of filters of A, and let $B = \bigcup_{i\in I} B_i$. For any $x, y \in B$, there are $B_i, B_j \in \mathcal{B}$ with $x \in B_i$ and $y \in B_j$. Since $\mathcal{B}$ is directed, there is $B_k \in \mathcal{B}$ with $B_i, B_j \subseteq B_k$. It follows that $x, y, x \wedge y \in B_k$ and hence $x \wedge y \in B$ also. For any $x \in B$, there is $B_i \in \mathcal{B}$ with $x \in B_i$, and for any $y \in A$ with $x \leq y$, $y \in B_i$ also, and consequently $y \in B$. Hence B is also a filter. Trivially, $A \in \mathfrak{F}(A)$, and it follows from Definition 3.22 that $\mathfrak{F}(A)$ is a topped algebraic $\bigcap$-structure on A. □

We shall use the set $\mathfrak{F}(\mathfrak{T}_\mathcal{D})$ of filters of $\mathfrak{T}_\mathcal{D}$ to define a closure operator on $\mathfrak{T}_\mathcal{D}$ as follows.

Definition 3.27. Let $C_F : \mathcal{P}(\mathfrak{T}_\mathcal{D}) \longrightarrow \mathcal{P}(\mathfrak{T}_\mathcal{D}) : U \mapsto \bigcap\{V \in \mathfrak{F}(\mathfrak{T}_\mathcal{D}) \mid U \subseteq V\}$. We define the map $C_T : \mathcal{P}(\mathfrak{T}_\mathcal{D}) \longrightarrow \mathcal{P}(\mathfrak{T}_\mathcal{D})$ to be such that for any $U \subseteq \mathfrak{T}_\mathcal{D}$,

$$C_T(U) \;=\; \begin{cases} \mathfrak{T}_\mathcal{D} & \text{if for some } t \in \mathfrak{T}_\mathcal{D},\, t \text{ and } t' \text{ are both in } C_F(U) \\ C_F(U) & \text{otherwise} \end{cases}$$

Proposition 3.28. *The map C_T is an algebraic closure operator on $\mathfrak{T}_\mathcal{D}$.*

Proof. We show first that C_T satisfies Conditions **CL1–CL3** of Definition 3.3.

(**CL1**). For any $U \subseteq \mathfrak{T}_\mathcal{D}$, if $t, t' \in C_F(U)$ for some $t \in \mathfrak{T}_\mathcal{D}$ then $C_T(U) = \mathfrak{T}_\mathcal{D}$ and $U \subseteq C_T(U)$. Otherwise, $C_T(U) = C_F(U)$ and by Proposition 3.26, $U \subseteq C_F(U)$.

(**CL2**). For any $U, V \subseteq \mathfrak{T}_\mathcal{D}$ with $U \subseteq V$, if $t, t' \in C_F(V)$ for some $t \in \mathfrak{T}_\mathcal{D}$, then $C_T(V) = \mathfrak{T}_\mathcal{D}$ and trivially $C_T(U) \subseteq C_T(V)$. If $t, t' \in C_F(U)$ for some $t \in \mathfrak{T}_\mathcal{D}$, then since $C_F(U) \subseteq C_F(V)$ by Proposition 3.26, it follows that $C_T(U) = C_T(V) = \mathfrak{T}_\mathcal{D}$ and again $C_T(U) \subseteq C_T(V)$. Otherwise, by Proposition 3.26 we have $C_T(U) = C_F(U) \subseteq C_F(V) = C_T(V)$.

(**CL3**). For any $U \subseteq \mathfrak{T}_\mathcal{D}$, if $t, t' \in C_F(U)$ for some $t \in \mathfrak{T}_\mathcal{D}$, $C_T(C_T(U)) = C_T(\mathfrak{T}_\mathcal{D}) = \mathfrak{T}_\mathcal{D} = C_T(U)$. If not, $C_T(C_T(U)) = C_T(C_F(U)) = C_F(C_F(U)) = C_F(U) = C_T(U)$, by Proposition 3.26.

Hence C_T is a closure operator on $\mathfrak{T}_{\mathcal{D}}$. Let $\mathcal{T} = \{T_i\}_{i\in I}$ be a directed family of subsets of $\mathfrak{T}_{\mathcal{D}}$. For any $i \in I$, either $C_T(T_i) = \mathfrak{T}_{\mathcal{D}}$ or $C_T(T_i) = C_F(T_i)$, and hence either $\bigcup_{i\in I} C_T(T_i) = \mathfrak{T}_{\mathcal{D}}$ or $\bigcup_{i\in I} C_T(T_i) = \bigcup_{i\in I} C_F(T_i)$. By Proposition 3.26 and Theorem 3.24, $\bigcup_{i\in I} C_F(T_i) = C_F(\bigcup_{i\in I} T_i)$, which is a closed set, and it follows that $\bigcup_{i\in I} C_T(T_i)$ is also a closed set. Trivially,

$$\begin{array}{lcccl} \forall i \in I.[T_i \subseteq \bigcup_{i\in I} T_i] & \Rightarrow & \forall i \in I.[C_T(T_i) \subseteq C_T(\bigcup_{i\in I} T_i)] & \Rightarrow & \bigcup_{i\in I} C_T(T_i) \subseteq C_T(\bigcup_{i\in I} T_i) \text{ and also} \\ \forall i \in I.[T_i \subseteq C_T(T_i)] & \Rightarrow & \bigcup_{i\in I} T_i \subseteq \bigcup_{i\in I} C_T(T_i) & \Rightarrow & C_T(\bigcup_{i\in I} T_i) \subseteq C_T(\bigcup_{i\in I} C_T(T_i)) \end{array}$$

and because $C_T(\bigcup_{i\in I} C_T(T_i)) = \bigcup_{i\in I} C_T(T_i)$, we then have $C_T(\bigcup_{i\in I} T_i) = \bigcup_{i\in I} C_T(T_i)$. By Theorem 3.24, C_T is then an algebraic closure operator on $\mathfrak{T}_{\mathcal{D}}$. □

From Definition 3.21, $(\mathbf{T}, C_T)$ is a logic. The next result is then a useful consequence of Definition 3.27.

Proposition 3.29. *For any $U \subseteq \mathfrak{T}_{\mathcal{D}}$, $C_T(U) = \mathfrak{T}_{\mathcal{D}}$ if and only if $\prod U \leftrightarrow \mathbf{0}$.*

Proof. (⇒). From Definition 3.27, for any $U \subseteq \mathfrak{T}_{\mathcal{D}}$, $C_F(U) = \uparrow(\prod U) = C_F(\{\prod U\})$ since $\mathfrak{T}_{\mathcal{D}}$ is a complete lattice. If $C_T(U) = \mathfrak{T}_{\mathcal{D}}$, we must then have $t, t' \in \uparrow(\prod U)$ for some $t \in \mathfrak{T}_{\mathcal{D}}$. It follows that $C_T(\{\prod U\}) = \mathfrak{T}_{\mathcal{D}} = C_T(\{\mathbf{0}\})$, and hence $\prod U \leftrightarrow \mathbf{0}$.

(⇐). If $\prod U \leftrightarrow \mathbf{0}$, then $C_T(\{\prod U\}) = \mathfrak{T}_{\mathcal{D}}$, and so $t, t' \in C_F(\{\prod U\})$ for some $t \in \mathfrak{T}_{\mathcal{D}}$. It follows that $t, t' \in C_F(U)$ also, and hence $C_T(U) = \mathfrak{T}_{\mathcal{D}}$. □

Let $\mathbb{V} = [X_0 \longrightarrow \mathfrak{T}_{\mathcal{D}}]$. We now define our putative consequence function C_n as follows.

Definition 3.30. We define the set-to-set function $C_n : \mathcal{P}(X) \longrightarrow \mathcal{P}(X)$ to be such that for any $U \subseteq X$,

$$C_n(U) \quad = \quad \{x \in X \mid \forall u \in \mathbb{V}.[u(x) \in C_T(u(U))]\}$$

The next result is a useful consequence of Definitions 3.27 and 3.30.

Lemma 3.31. *For any $U \subseteq X$ and any $u \in \mathbb{V}$, $C_T(u(C_n(U))) \subseteq C_T(u(U))$.*

Proof. Let U be a subset of X. For any $u \in \mathbb{V}$ we have

$$\begin{array}{rcll} x \in C_n(U) & \Rightarrow & u(x) \in C_T(u(U)) & \text{(Definition 3.30)} \\ & \Rightarrow & u(C_n(U)) \subseteq C_T(u(U)) & (u(C_n(U)) = \{u(x) \mid x \in C_n(U)\}) \\ & \Rightarrow & C_T(u(C_n(U))) \subseteq C_T(C_T(u(U))) = C_T(u(U)) & \text{(since } C_T \text{ is a closure operator)} \end{array}$$

which gives us the result, as required. □

Proposition 3.32. *The set-to-set function C_n is an algebraic closure operator on X.*

Proof. We show first that C_n satisfies Conditions **CL1–CL3** of Definition 3.3.

(**CL1**). For any $x \in U$ and any $u \in \mathbb{V}$, $u(x) \in u(U)$, and hence $u(x) \in C_T(u(U))$. It follows that for any $U \subseteq X$, $U \subseteq C_n(U)$.

(**CL2**). For $U, V \subseteq X$, suppose that $U \subseteq V$. Then

$$\begin{aligned} x \in C_n(U) \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(u(U))] \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(u(V))] && \text{(property of } C_T\text{)} \\ \Rightarrow\ & x \in C_n(V) \end{aligned}$$

so that $C_n(U) \subseteq C_n(V)$.

(**CL3**). By **CL1**, $C_n(U) \subseteq C_n(C_n(U))$. For any $x \in X$,

$$\begin{aligned} x \in C_n(C_n(U)) \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(u(C_n(U)))] \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(u(U))] && \text{(Lemma 3.31)} \\ \Rightarrow\ & x \in C_n(U) \end{aligned}$$

Thus $C_n(C_n(U)) \subseteq C_n(U)$, so that $C_n(U) = C_n(C_n(U))$.

Hence C_n is a closure operator on X. Now let $\mathcal{X} = \{X_i\}_{i \in I}$ be a directed family of subsets of X. For any $x \in X$,

$$\begin{aligned} & x \in C_n(\bigcup_{i \in I} C_n(X_i)) \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(u(\bigcup_{i \in I} C_n(X_i)))] \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in C_T(\bigcup_{i \in I} u(C_n(X_i)))] \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in \bigcup_{i \in I} C_T(u(C_n(X_i)))] && \text{(since } \{u(C_n(X_i))\}_{i \in I} \text{ is directed and } C_T \text{ is algebraic)} \\ \Rightarrow\ & \forall u \in \mathbb{V}.[u(x) \in \bigcup_{i \in I} C_T(u(X_i))] && \text{(Lemma 3.31)} \\ \Rightarrow\ & x \in \bigcup_{i \in I} C_n(X_i) && \text{(Definition 3.30)} \end{aligned}$$

and hence $C_n(\bigcup_{i \in I} C_n(X_i)) \subseteq \bigcup_{i \in I} C_n(X_i)$. By **CL1** of Definition 3.3, $\bigcup_{i \in I} C_n(X_i) \subseteq C_n(\bigcup_{i \in I} C_n(X_i))$, and it follows that $C_n(\bigcup_{i \in I} C_n(X_i)) = \bigcup_{i \in I} C_n(X_i)$, so that $\bigcup_{i \in I} C_n(X_i)$ is a closed set. Trivially,

$$\begin{aligned} \forall i \in I.[X_i \subseteq \bigcup_{i \in I} X_i] \ &\Rightarrow\ \forall i \in I.[C_n(X_i) \subseteq C_n(\textstyle\bigcup_{i \in I} X_i)] \ \Rightarrow\ \bigcup_{i \in I} C_n(X_i) \subseteq C_n(\bigcup_{i \in I} X_i) \quad \text{and also} \\ \forall i \in I.[X_i \subseteq C_n(X_i)] \ &\Rightarrow\ \textstyle\bigcup_{i \in I} X_i \subseteq \bigcup_{i \in I} C_n(X_i) \ \Rightarrow\ C_n(\displaystyle\bigcup_{i \in I} X_i) \subseteq C_n(\bigcup_{i \in I} C_n(X_i)) \end{aligned}$$

so that $C_n(\bigcup_{i \in I} X_i) \subseteq C_n(\bigcup_{i \in I} C_n(X_i)) = \bigcup_{i \in I} C_n(X_i)$, and hence $C_n(\bigcup_{i \in I} X_i) = \bigcup_{i \in I} C_n(X_i)$. By Theorem 3.24, C_n is then an algebraic closure operator on X. □

Because C_n is a closure operator on X, we may use it as the consequence function of our logic $\mathfrak{L}_X$.

Proposition 3.33. *For any $U \subseteq X$, $C_n(U) = X$ if and only if for every $u \in \mathbb{V}$, $C_T(u(U)) = \mathfrak{T}_{\mathcal{D}}$.*

Proof. (⇒). Let $U \subseteq X$ be such that $C_n(U) = X$. Then for any $x \in X$, $x, x' \in C_n(U)$. But then

$$\begin{aligned} x, x' \in C_n(U) \Rightarrow\ & \forall u \in \mathbb{V}.[u(x), u(x') \in C_T(u(U))] && \text{(Definition 3.30)} \\ \Rightarrow\ & \forall u \in \mathbb{V}.\exists t \in \mathfrak{T}_{\mathcal{D}}.[t, t' \in C_T(u(U))] && \text{(use } t = u(x)\text{)} \\ \Rightarrow\ & \forall u \in \mathbb{V}.[C_T(u(U)) = \mathfrak{T}_{\mathcal{D}}] && \text{(Definition 3.27)} \end{aligned}$$

(⇐). This follows directly from the definitions of C_T and C_n. □

From Definition 2.4(iv) in [12], given a logic $\mathfrak{L}_A = (\mathbf{A}, C_n)$, two sentences $x, y \in A$ are logically equivalent, written $x \leftrightarrow y$, if each is a logical consequence of the other. From Definition 3.21, we then have the following result (compare also Example 3.9(iv) in [12]).

Proposition 3.34. *Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic. For any $x, y \in A$,*

$$x \leftrightarrow y \quad \textit{if and only if} \quad x \in C_n(\{y\}) \textit{ and } y \in C_n(\{x\})$$

Proof. Let $x, y \in A$. If $x \leftrightarrow y$ then $C_n(\{x\}) = C_n(\{y\})$. By Condition **L1** of Definition 3.21 we have $x \in C_n(\{x\})$ and $y \in C_n(\{y\})$, so $x \in C_n(\{y\})$ and also $y \in C_n(\{x\})$. For the reverse direction,

$$x \in C_n(\{y\}) \quad \Rightarrow \quad C_n(\{x\}) \subseteq C_n(C_n(\{y\})) \subseteq C_n(\{y\})$$

by Condition **L2** of Definition 3.21. Similarly, interchanging x and y gives $C_n(\{y\}) \subseteq C_n(\{x\})$, from which it follows that $C_n(\{x\}) = C_n(\{y\})$, and hence that $x \leftrightarrow y$. □

Within $\mathfrak{L}_X = (\mathbf{X}, C_n)$, the consequence function C_n can thus be thought of as imposing a logical structure on the given set X of sentences. With respect to this structure, if for any $x, y \in X$ we have $x \leftrightarrow y$, then x does not convey any more information than y, and so we may think of x and y as representing the same sentence in X. Consequently, we may partition X by logical equivalence, taking $x \approx y$ if and only if $x \leftrightarrow y$, and henceforth we assume that X contains only sentences that are logically distinct.

At this point we can do little with $\mathfrak{L}_X$ since we have yet to specify any logical operations that we can carry out within it. In the sub-sections to follow, we remedy this omission, beginning with negation and concluding with conjunction and disjunction. As part of our treatment of negation we consider the related aspect of contradiction and we discuss how one sentence might be compatible with a given set of sentences.

3.2.2 Negation and Contradiction

Contradiction plays an important role in the defining the idea of negation within a logic. From [12], it is characteristic of contradiction in logics such as classical and intuitionistic logic that any proposition follows from a contradiction.

We first observe the following result.

Proposition 3.35. *Within $\mathbf{X} = (X, +, *, ')$, the operation $'$ is an involution on X.*

Proof. For any $x, y \in X$,

$$\begin{aligned} x \leq_0 y \Rightarrow \quad & \forall U \in \Omega X \mid x, y \in U.[\forall u \in W_{x,U} \cup W_{y,U}.[u(x) \leq_T u(y)]] \quad \text{(from Definition 2.31)} \\ \Rightarrow \quad & \forall U \in \Omega X \mid x, y \in U.[\forall u \in W_{x,U} \cup W_{y,U}.[u(y') \leq_T u(x')]] \quad \text{(property of involution)} \\ \Rightarrow \quad & y' \leq_0 x' \end{aligned}$$

From Definition 2.31,

$$\begin{aligned} x \leq_X y \quad & \Rightarrow \quad \exists x_1, x_2, \ldots, x_n \in X.[x \leq_0 x_1 \leq_0 x_2 \leq_0 \ldots \leq_0 x_n \leq_0 y] \\ & \Rightarrow \quad \exists x_1, x_2, \ldots, x_n \in X.[y' \leq_0 x_n' \leq_0 x_{n-1}' \leq_0 \ldots \leq_0 x_1' \leq_0 x'] \\ & \Rightarrow \quad y' \leq_X x' \end{aligned}$$

Because $'$ was defined such that $x'' = x$, by Definition 3.18 $'$ is then an involution. □

From p115 in [12], we have the following.

Definition 3.36. Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic of type $\mathcal{F}$. An element $\mathbf{0}_A \in A$ is a *contradiction* if $C_n(\{\mathbf{0}_A\}) = A$. An element $\mathbf{1}_A \in A$ is called *theoremhood* if for any $x \in A$, $\mathbf{1}_A \in C_n(\{x\})$.

To cater for theoremhood and contradiction in $\mathfrak{L}_X = (\mathbf{X}, C_n)$, we extend the set X_0 of primitives to contain the elements $\mathbf{0}_X$ and $\mathbf{1}_X$. For any $w \in \mathbb{V}$ we define $w(\mathbf{0}_X) = \mathbf{0}$ and $w(\mathbf{1}_X) = \mathbf{1}$.

Lemma 3.37. *For any $x \in X$, $\mathbf{0}_X \leq_X x \leq_X \mathbf{1}_X$.*

Proof. This follows directly from Definition 2.31 by observing that $\mathbf{T} = (\mathfrak{T}_{\mathcal{D}}, +, *, ', \mathbf{0}, \mathbf{1}, \leq_T)$ is a bounded r-algebra. □

Observe that from Lemma 3.37, $\mathbf{0}'_X = \mathbf{1}_X$ and $\mathbf{1}'_X = \mathbf{0}_X$.

Proposition 3.38. *Within $\mathfrak{L}_X = (\mathbf{X}, C_n)$, $\mathbf{0}_X$ is a contradiction and $\mathbf{1}_X$ is theoremhood.*

Proof. For any $u \in \mathbb{V}$, $u(\mathbf{0}_X) = \mathbf{0}$, and hence $C_T(\{u(\mathbf{0}_X)\}) = \mathfrak{T}_{\mathcal{D}}$. By Proposition 3.33, $C_n(\{\mathbf{0}_X\}) = X$ and hence $\mathbf{0}_X$ is a contradiction in $\mathfrak{L}_X$. Similarly, for any $x \in X$ we have

$$\begin{array}{lll} & \forall u \in \mathbb{V}.[\mathbf{1} \in C_T(\{u(x)\})] & \text{(Definitions 3.25 and 3.27; } \forall t \in \mathfrak{T}_{\mathcal{D}}.[t \leq_T \mathbf{1}]) \\ \Rightarrow & \forall u \in \mathbb{V}.[u(\mathbf{1}_X) \in C_T(\{u(x)\})] & \\ \Rightarrow & \mathbf{1}_X \in C_n(\{x\}) & \text{(Definition 3.30)} \end{array}$$

and hence $\mathbf{1}_X$ is theoremhood in $\mathfrak{L}_X$. □

For any $x \in X$, if $C_n(\{x\}) = X$ then $x \leftrightarrow \mathbf{0}_X$ (compare Definition 3.21 and Proposition 3.34), and we may replace any occurrence of x in a sentence or set of sentences by $\mathbf{0}_X$. We will use this observation without further comment.

Proposition 3.39 may be compared with Theorem 3.9 in [12].

Proposition 3.39. *For $\mathfrak{L}_X = (\mathbf{X}, C_n)$, let*

$$\begin{array}{lll} \mathcal{C}_X & = & \{C_n(U) \mid U \subseteq X\}, \textit{ the closed sets of } X \\ T_X & = & \bigcap\limits_{x \in X} C_n(\{x\}) \end{array}$$

Then,

i) Every closed set in X is non-empty.

ii) $T_X = \bigcap_{U \in \mathcal{C}_X} U$

iii) $Th(\mathfrak{L}_X) = T_X$

iv) $C_n(\{\mathbf{1}_X\}) = C_n(\emptyset)$

Proof. (i). By **CL2** of Definition 3.3, for any $U \subseteq X$ and any $x \in U$, $C_n(\{x\}) \subseteq C_n(U)$. By Proposition 3.38 and Definition 3.36, for any $x \in X$, $\mathbf{1}_X \in C_n(\{x\})$, and it follows then that $C_n(U) \neq \emptyset$.

(ii). Trivially, $\{C_n(\{x\})\}_{x \in X} \subseteq \mathcal{C}_X$, so $\bigcap_{U \in \mathcal{C}_X} U \subseteq \bigcap_{x \in X} C_n(\{x\}) = T_X$ by Lemma 2.22(v) in [13]. By (i), for any $U \in \mathcal{C}_X$, there is $x \in U$ and hence $T_X \subseteq C_n(\{x\}) \subseteq C_n(U)$. It follows that $T_X \subseteq \bigcap_{U \in \mathcal{C}_X} U$,

and hence $T_X = \bigcap_{U \in \mathcal{C}_X} U$.
(iii). From Proposition 3.4, for any $U \subseteq X$, $C_n(U) = \bigcap\{V \in \mathcal{C}_X \mid U \subseteq V\}$. In particular,

$$Th(\mathfrak{L}_X) = C_n(\emptyset) = \bigcap\{V \in \mathcal{C}_X \mid \emptyset \subseteq V\} = \bigcap_{U \in \mathcal{C}_X} U = T_X$$

(iv). Trivially, $C_n(\{\mathbf{1}_X\}) = \{\mathbf{1}_X\}$ and by (i), for any $U \in \mathcal{C}_X$, $\{\mathbf{1}_X\} \subseteq U$. It now follows that $C_n(\emptyset) = T_X = \{\mathbf{1}_X\} = C_n(\{\mathbf{1}_X\})$. □

In particular, Proposition 3.39(iv) means that $\mathfrak{L}_X$ has a non-empty set of theorems.

We would like to use the n-ary fundamental operations of (X, F), where $n \geq 1$, as logical operations. From Definition 3.11 in [12], we have the following.

Definition 3.40. Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic of type $\mathcal{F}$, and for $n \geq 1$ let $f \in \mathcal{F}$ be an n-ary function symbol. The fundamental operation $f^{\mathbf{A}}$ of $\mathbf{A}$ is called a *logical operation* or *connective* in $\mathfrak{L}_A$ if it preserves logical equivalence. That is, for any $a_1, a_2, \ldots, a_n, b_1, b_2, \ldots, b_n \in A$, if for all $1 \leq i \leq n$ we have that a_i is logically equivalent to b_i, then $f^{\mathbf{A}}(a_1, a_2, \ldots, a_n)$ is logically equivalent to $f^{\mathbf{A}}(b_1, b_2, \ldots, b_n)$.

Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic with a unary operation ν. For any $x \in A$, we would like $\nu(x)$ to be the negation of x. Intuitively, for this to be the case, ν must satisfy certain properties. Negation is related to contradiction by the law of contradiction, which from p115 in [12] may be stated as

N1: For any $x \in A$, $C_n(\{x, \nu(x)\}) = A$

Axiom **N1** is too weak for ν to preserve logical equivalence, as required by Definition 3.40. To **N1** we therefore add the following axioms, from p116 in [12]:

N3.1: For any $x, y \in A$ and $X \subseteq A$, $x \in C_n(X \cup \{y\})$ implies $\nu(y) \in C_n(X \cup \{\nu(x)\})$

N3.2: For any $x, y \in A$ and $X \subseteq A$, $\nu(x) \in C_n(X \cup \{y\})$ implies $\nu(y) \in C_n(X \cup \{x\})$

N3.3: For any $x, y \in A$ and $X \subseteq A$, $x \in C_n(X \cup \{\nu(y)\})$ implies $y \in C_n(X \cup \{\nu(x)\})$

The versions of these axioms in which the set variable X is suppressed (effectively by setting X to $\emptyset$) are denoted by $\mathbf{N3.1}^0$, $\mathbf{N3.2}^0$ and $\mathbf{N3.3}^0$. Only Axiom **N3.1** is needed for ν to be a logical operation, as shown by the following result, based on Theorem 4.6 in [12]. In the proof of Theorem 3.41, we have suppressed the set variable from the notation, so that ν satisfies $\mathbf{N3.1}^0$ rather than **N3.1**. This suppression is permissible because we are proving that ν is a logical rather than a *strongly* logical operation (compare Definition 3.11 in [12]).

Theorem 3.41. *Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic with a unary operation ν that satisfies Axiom* **N3.1***. Then ν is a logical operation.*

Proof. For $x, y \in A$ we have

$$\begin{aligned} x \leftrightarrow y \quad &\Rightarrow \quad x \in C_n(\{y\}) \text{ and } y \in C_n(\{x\}) \\ &\Rightarrow \quad \nu(y) \in C_n(\{\nu(x)\}) \text{ and } \nu(x) \in C_n(\{\nu(y)\}) && \text{(Axiom } \mathbf{N3.1} \text{ with } X = \emptyset) \\ &\Rightarrow \quad \nu(x) \leftrightarrow \nu(y) && \text{(Proposition 3.34)} \end{aligned}$$

and so by Definition 3.40, ν is a logical operation. □

Proposition 3.42 and Lemma 3.43 show that $'$ may indeed serve as a negation within $(X, +, *, ')$.

Proposition 3.42. *Within $(X, +, *, ')$, the operation $'$ satisfies Axiom* **N1**.

Proof. This follows directly from Definitions 3.27 and 3.30. □

As in the proof of Theorem 3.41, we suppress the set variable in the proof of the next result.

Lemma 3.43. *Within $(X, +, *, ')$, the operation $'$ satisfies Axioms* $\mathbf{N3.1}^0$, $\mathbf{N3.2}^0$ *and* $\mathbf{N3.3}^0$.

Proof. For any $x, y \in X$,

$$\begin{array}{lll} & x \in C_n(\{y\}) & \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x) \in C_T(\{u(y)\})] & \\ \Rightarrow & \forall u \in \mathbb{V}.[u(y) \leq_T u(x)] & (\text{since } u(y) \leftrightarrow \mathbf{0} \text{ or } C_T(\{u(y)\}) = C_F(\{u(y)\}) = \uparrow u(y)) \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x') \leq_T u(y')] & (u(x') = (u(x))'; \text{ property of } ') \\ \Rightarrow & \forall u \in \mathbb{V}.[u(y') \in C_T(\{u(x')\})] & (\text{since } u(x') \leftrightarrow \mathbf{0} \text{ or } C_T(\{u(x')\}) = C_F(\{u(x')\}) = \uparrow u(x')) \\ \Rightarrow & y' \in C_n(\{x'\}) & (\text{definition of } C_n) \end{array}$$

and hence $'$ satisfies $\mathbf{N3.1}^0$. Similarly, by substituting x' for x, we can show that $'$ satisfies $\mathbf{N3.2}^0$, and by substituting y' for y, we can show that $'$ satisfies $\mathbf{N3.3}^0$. □

We may therefore extend our logic to include the involution $'$, so that $\mathfrak{L}_X = ((X, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$, where $\mathbf{0}_X$ represents contradiction and $\mathbf{1}_X$ represents theoremhood in $\mathfrak{L}_X$.

The next result shows that for $\mathfrak{L}_X$, for any $U \subseteq X$ it is the case that $C_n(U) = X$ if and only if either $\mathbf{0}_X \in U$ or there is $x \in U$ with $x \leftrightarrow \mathbf{0}_X$.

Proposition 3.44. *For any $U \subseteq X$,*

$$\exists x \in U.[x \leftrightarrow \mathbf{0}_X] \text{ if and only if } C_n(U) = X$$

Proof. ($\Rightarrow$). Let $U \subseteq X$ be such that there is $x \in U$ with $x \leftrightarrow \mathbf{0}_X$. Then since $C_n(\{x\}) \subseteq C_n(U)$ and $C_n(\{x\}) = C_n(\{\mathbf{0}_X\}) = X$, it follows that $C_n(U) = X$.

($\Leftarrow$). Let $U \subseteq X$ be such that $C_n(U) \neq X$. Then, in particular, for no $x \in U$ is it the case that $C_n(\{x\}) = X$, and hence there does not exist $x \in U$ with $x \leftrightarrow \mathbf{0}_X$. □

From C_n, we have the notion that a sentence x is compatible with a given set $U \subseteq X$ of sentences. Definition 3.45 may be compared with Definition 3.12 in [36].

Definition 3.45. Let $x \in X$ be a sentence. A set $U \subseteq X$ is *x-compatible* if $x' \notin C_n(U)$. It is a *maximal x-compatible set* if it is x-compatible and for any other $V \subseteq X$ with $U \subset V$, $x' \in C_n(V)$.

Proposition 3.46 is based on Theorem 4.8 in [12] and makes the role of x-compatibility clearer. The proof is given in [12], so we will not restate it here.

Proposition 3.46. *Let $\mathfrak{L}_X = ((X, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$ be a logic in which $\mathbf{0}_X$ is a contradiction, $\mathbf{1}_X$ is theoremhood and $'$ satisfies Axiom* **N1**. *Then $\mathfrak{L}_X$ satisfies*

N4: *For any $x \in X$ and $U \subseteq X$, $\mathbf{0}_X \in C_n(U \cup \{x\})$ if and only if $x' \in C_n(U)$*

if and only if ′ satisfies Axioms **N3.1** *and* **N3.2**.

Now suppose that $V_0 \subseteq U$ is an x-compatible set of sentences, for $x \in X$. We need $V \subseteq U$ such that $V_0 \subseteq V$ and V is a maximal x-compatible set. To show that such an element exists, we appeal to Zorn's Lemma in the form of the following statement (see Definition 10.2 in [13]):

> *Let P be a non-empty ordered set in which every non-empty chain has an upper bound. Then P has a maximal element.*

This statement may specialised to families of sets as follows:

> *Let $\mathcal{A}$ be a non-empty family of sets such that $\bigcup_{i\in I} A_i \in \mathcal{A}$ whenever $\{A_i\}_{i\in I}$ is a non-empty chain in $(\mathcal{A}, \subseteq)$. Then $\mathcal{A}$ has a maximal element.*

Let $\mathcal{V} = \{V \subseteq U \mid V_0 \subseteq V$ and V is x-compatible$\}$. Then $V_0 \in \mathcal{V}$ so $\mathcal{V} \neq \emptyset$. Let $\mathcal{V}_c = \{V_i\}_{i\in I}$ be a non-empty chain in $\mathcal{V}$, and let $V_c = \bigcup_{i\in I} V_i$. Since $\mathcal{V}_c$ is directed, $C_n(V_c) = C_n(\bigcup_{i\in I} V_i) = \bigcup_{i\in I} C_n(V_i)$ because C_n is algebraic. Since each V_i is x-compatible, so is V_c and hence $V_c \in \mathcal{V}$. Since $\mathcal{V}_c$ was arbitrary, it follows that any non-empty chain in $\mathcal{V}$ has an upper bound in $\mathcal{V}$ and by Zorn's Lemma, $\mathcal{V}$ then contains a maximal element. This establishes the following result (compare Lemma 3.13 in [36]).

Proposition 3.47. *For the logic $\mathfrak{L}_X = (\mathbf{X}, C_n)$, let $U \subseteq X$ be a set of sentences, let $x \in X$ be a sentence and let $V \subseteq U$ be an x-compatible set of sentences. Then V is contained in a maximal x-compatible subset of U.*

Remark 3.48. In Definition 3.12 and Lemma 3.13 of [36], the authors describe x-compatibility in terms of fallbacks, and hence in terms of deductively closed sets. Definition 3.45 and Proposition 3.47 are effectively expressed in terms of belief bases.

In the next section, we set out the formal requirements for conjunction and disjunction as we have just done for negation.

3.2.3 Conjunction and Disjunction

As a notational convention, let $\mathcal{P}_f(Y)$ denote the family of finite subsets of a set Y. Definition 3.49 is adapted from the axioms presented on p118 in [12].

Definition 3.49. Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic of type $\mathcal{F}$. We take conjunction π and disjunction σ in $\mathfrak{L}_A$ to be functions such that for each non-empty subset $U \in \mathcal{P}_f(A)$, $\pi(U)$ and $\sigma(U)$ are sentences in A. For any $S, T \in \mathcal{P}(A)$ and any non-empty $U \in \mathcal{P}_f(A)$, π and σ have the properties

Cn1: For any $x \in A$, $x \in U$ implies that $x \in C_n(S \cup \{\pi(U)\})$

Cn2: If for all $x \in U$ we have $x \in C_n(S \cup T)$, then also $\pi(U) \in C_n(S \cup T)$

Dn1: For any $x \in A$, $x \in U$ implies that $\sigma(U) \in C_n(S \cup \{x\})$

Dn2: For any $y \in A$, if y is such that for all $x \in U$, $y \in C_n(S \cup \{x\})$, then $y \in C_n(S \cup \{\sigma(U)\})$

An operation that satisfies **Cn1** and **Cn2** is called a *normal conjunction (on A)*. Similarly, an operation that satisfies **Dn1** and **Dn2** is called a *normal disjunction (on A)*.

As defined, π and σ operate on finite sets. We would like to replace them with the binary operations $*$ and $+$ in the algebra $\mathbf{X}$ of sentences of our logic $\mathfrak{L}_X$. To do this, we first need the following results. Lemma 3.50 is adapted from Theorem 3.7 in [12], while Lemma 3.51 and Corollary 3.52 are adapted from Corollaries 4.10 and 4.11 in [12]. The proofs of Lemmas 3.50 and 3.51 follow by definition-chasing, and Corollary 3.52 then follows by application of Lemmas 3.50 and 3.51, so we will omit the proofs here.

Lemma 3.50. *Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic. Then for any $S, T \subseteq A$,*

$$C_n(S \cup T) \quad = \quad C_n(S \cup C_n(T)) \quad = \quad C_n(C_n(S) \cup C_n(T))$$

Lemma 3.51. *Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic with normal conjunction π and normal disjunction σ. Then for all $S, T \in \mathcal{P}(A)$ and non-empty $U \in \mathcal{P}_f(A)$,*

i) $C_n(\{\pi(U)\}) = C_n(U)$

ii) $C_n(S \cup \{\sigma(U)\}) = \bigcap_{x \in U} C_n(S \cup \{x\})$

Corollary 3.52. *Let $\mathfrak{L}_A = (\mathbf{A}, C_n)$ be a logic with normal conjunction π and normal disjunction σ. Then for all non-empty $S, T \in \mathcal{P}_f(A)$,*

i) $\pi(S \cup T)$ *and* $\pi(\{\pi(S), \pi(T)\})$ *are logically equivalent.*

ii) $\sigma(S \cup T)$ *and* $\sigma(\{\sigma(S), \sigma(T)\})$ *are logically equivalent.*

For any non-empty finite subset $U = \{x_i\}_{i \in I}$ of X, let us define

$$\prod U \quad \triangleq \quad x_1 * x_2 * \ldots * x_n$$
$$\sum U \quad \triangleq \quad x_1 + x_2 + \ldots + x_n$$

Proposition 3.53. *Within the logic $\mathfrak{L}_X = (\mathbf{X}, C_n)$, the operations $\prod$ and $\sum$ are respectively a normal conjunction and a normal disjunction on X.*

Proof. Let $U = \{x_i\}_{i \in I}$ be a non-empty finite subset of X. For any $x \in U$ we have

$$\begin{array}{lll} & \forall u \in \mathbb{V}.[u(\prod U) \leq_T u(x)] & \text{(since } u(\prod U) = u(x_1) * u(x_2) * \ldots * u(x_n)\text{)} \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x) \in C_T(\{u(\prod U)\})] & \text{(property of } C_T\text{)} \\ \Rightarrow & x \in C_n(\{\prod U\}) & \text{(Definition 3.30)} \\ \Rightarrow & x \in C_n(S \cup \{\prod U\}) & \text{(property \textbf{CL2} of Definition 3.3)} \end{array}$$

and hence $\prod$ satisfies **Cn1**. Similarly, for any $S, T \subseteq X$,

$$\begin{array}{lll} & \forall x \in U.[x \in C_n(S \cup T)] & \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x_1), \ldots, u(x_n) \in C_T(\{u(S \cup T)\})] & \text{(Definition 3.30)} \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x_1) * \ldots * u(x_n) \in C_T(\{u(S \cup T)\})] & \text{(property of } C_T\text{)} \\ \Rightarrow & \forall u \in \mathbb{V}.[u(\prod U) \in C_T(\{u(S \cup T)\})] & \\ \Rightarrow & \prod U \in C_n(S \cup T) & \text{(Definition 3.30)} \end{array}$$

and hence $\prod$ satisfies **Cn2**. For any $S \subseteq X$ and $x \in U$,

$$\forall u \in \mathbb{V}.[u(x) \leq_T u(\sum U)] \qquad \text{(since } u(\sum U) = u(x_1) + \ldots + u(x_n)\text{)}$$

$$
\begin{array}{lll}
\Rightarrow & \forall u \in \mathbb{V}.[u(\sum U) \in C_T(\{u(x)\})] & \text{(property of } C_T\text{)} \\
\Rightarrow & \sum U \in C_n(\{x\}) & \text{(Definition 3.30)} \\
\Rightarrow & \sum U \in C_n(S \cup \{x\}) & \text{(property \textbf{CL2} of Definition 3.3)}
\end{array}
$$

and hence $\sum$ satisfies **Dn1**. Finally, for any $y \in X$,

$$
\begin{array}{rll}
y \in \bigcap_{x \in U} C_n(S \cup \{x\}) \Rightarrow & \forall x \in U.\forall u \in \mathbb{V}.[u(y) \in C_T(u(S \cup \{x\}))] & \text{(Definition 3.30)} \\
\Rightarrow & \forall u \in \mathbb{V}.[u(y) \in \bigcap_{x \in U} C_T(u(S \cup \{x\}))] &
\end{array}
$$

For any $x \in U$ and $u \in \mathbb{V}$, $u(S \cup \{x\}) = u(S) \cup \{u(x)\}$. From Definition 3.27 we then have that $C_F(u(S \cup \{x\})) = C_F(u(S) \cup \{u(x)\})$. Denoting $\sum$ and $\prod$ in $\mathbf{T}$ respectively by $\bigvee$ and $\bigwedge$, and writing $t_S = \bigwedge u(S)$, $t_x = u(x)$ and $t_y = u(y)$, we then have that $C_F(u(S \cup \{x\})) = \uparrow(t_S * t_x)$. If now $C_T(u(S \cup \{x\})) = \mathfrak{T}_{\mathcal{D}}$, then $(t_S * t_x) \leftrightarrow \mathbf{0}$ by Proposition 3.29, and so $t_y \in \uparrow(t_S * t_x)$. Otherwise, $C_T(u(S \cup \{x\})) = \uparrow(t_S * t_x)$ and again $t_y \in \uparrow(t_S * t_x)$. Hence,

$$
\begin{array}{lll}
& \forall u \in \mathbb{V}.[u(y) \in \bigcap_{x \in U} C_T(u(S \cup \{x\}))] & \\
\Rightarrow & \forall u \in \mathbb{V}.[t_y \in \bigcap_{x \in U} \uparrow(t_S * t_x)] & \\
\Rightarrow & \forall u \in \mathbb{V}.\forall x \in U.[(t_S * t_x) \leq_T t_y] & \text{(property of up set)} \\
\Rightarrow & \forall u \in \mathbb{V}.[\bigvee_{x \in U} (t_S * t_x) \leq_T t_y] & \text{(property of least upper bound)} \\
\Rightarrow & \forall u \in \mathbb{V}.[t_S \wedge (\bigvee_{x \in U} t_x) \leq_T t_y] & \text{(since } \mathfrak{T}_{\mathcal{D}} \text{ is distributive, Proposition 3.17)} \\
\Rightarrow & \forall u \in \mathbb{V}.[t_y \in \uparrow(t_S \wedge (\bigvee_{x \in U} t_x))] & \text{(property of up set)} \\
\Rightarrow & \forall u \in \mathbb{V}.[u(y) \in C_T(u(S \cup \{\sum U\}))] & \text{(Definition 3.27)} \\
\Rightarrow & y \in C_n(S \cup \{\sum U\}) & \text{(Definition 3.30)}
\end{array}
$$

and hence $\sum$ satisfies **Dn2**. Taken together, it follows from Definition 3.49 that $\prod$ and $\sum$ are respectively normal conjunction and normal disjunction on X. □

Proposition 3.54. *The operations $*$ and $+$ are logical connectives in the logic $\mathfrak{L}_X = (\mathbf{X}, C_n)$.*

Proof. For any $x, y \in X$ we have $x * y = \prod\{x, y\}$, and so $C_n(\{x * y\}) = C_n(\{\prod\{x, y\}\})$. Suppose that $x \leftrightarrow a$ and $y \leftrightarrow b$, for $a, b \in X$. Then

$$
\begin{array}{rcll}
C_n(\{x * y\}) & = & C_n(\{\prod\{x, y\}\}) & \text{(since } x * y = \prod\{x, y\}\text{)} \\
& = & C_n(\{x, y\}) & \text{(by Lemma 3.51)} \\
& = & C_n(C_n(\{x\}) \cup C_n(\{y\})) & \text{(by Lemma 3.50)} \\
& = & C_n(C_n(\{a\}) \cup C_n(\{b\})) & \text{(since } x \leftrightarrow a \text{ and } y \leftrightarrow b\text{)} \\
& = & C_n(\{a, b\}) & \text{(by Lemma 3.50)} \\
& = & C_n(\{\prod\{a, b\}\}) & \text{(by Lemma 3.51)} \\
& = & C_n(\{a * b\}) &
\end{array}
$$

from which it follows that $*$ is a logical connective in $\mathfrak{L}_X$. Similarly, for any $S \subseteq X$,

$$
C_n(S \cup \{x + y\}) = C_n(S \cup \{\sum\{x, y\}\}) \qquad \text{(since } x + y = \sum\{x, y\}\text{)}
$$

$$\begin{aligned}
&= C_n(S \cup \{x\}) \cap C_n(S \cup \{y\}) && \text{(by Lemma 3.51)}\\
&= C_n(S \cup C_n(\{x\})) \cap C_n(S \cup C_n(\{y\})) && \text{(by Lemma 3.50)}\\
&= C_n(S \cup C_n(\{a\})) \cap C_n(S \cup C_n(\{b\})) && \text{(since } x \leftrightarrow a \text{ and } y \leftrightarrow b\text{)}\\
&= C_n(S \cup \{a\}) \cap C_n(S \cup \{b\}) && \text{(by Lemma 3.50)}\\
&= C_n(S \cup \{\textstyle\sum\{a,b\}\}) && \text{(by Lemma 3.51)}\\
&= C_n(S \cup \{a+b\})
\end{aligned}$$

and hence $+$ is a logical connective in $\mathfrak{L}_X$. □

We now have a logic $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$.

Recall from Section 3.2.1 that for $x, y \in X$, we write $x \approx y$ to mean $x \leftrightarrow y$. Proposition 3.55 may be compared to Theorem 4.15(iv) in [12].

Proposition 3.55. *Within the logic* $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$, $x * (y + z) \approx (x * y) + (x * z)$ *for any* $x, y, z \in X$.

Proof. For any $x, y, z \in X$,

$$\begin{aligned}
C_n(\{x * (y+z)\}) &= C_n(\{\textstyle\prod\{x, y+z\}\}) &&\\
&= C_n(\{x, y+z\}) && \text{(Lemma 3.51(i))}\\
&= C_n(\{x, y\}) \cap C_n(\{x, z\}) && \text{(Lemma 3.51(ii) with } S = \{x\}\text{)}\\
&= C_n(\{x * y\}) \cap C_n(\{x * z\}) && \text{(Lemma 3.51(i))}\\
&= C_n(\{(x * y) + (x * z)\}) && \text{(Lemma 3.51(ii))}
\end{aligned}$$

so that $x * (y+z)$ and $(x*y) + (x*z)$ are logically equivalent, and hence $x * (y+z) \approx (x*y) + (x*z)$ as required. □

Thus, given a sentence in conjunctive normal form (*i.e.* as a conjunction of disjuncts), we can expand the conjunctions as prescribed by the distributive law for lattices (see Definition 4.4(i) in [13]) and equivalently rewrite the sentence in disjunctive normal form (*i.e.* as a disjunction of conjuncts).

Proposition 3.56 captures the familiar absorption laws for lattices (see Theorem 2.9 in [13]; compare also Theorem 4.15(iv) in [12]).

Proposition 3.56. *Within the logic* $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$,

i) $x \approx x * (x + y)$

ii) $x \approx x + x * y$

for any $x, y \in X$.

Proof. Since $\mathfrak{T}_\mathcal{D}$ is a lattice, we may apply the absorption laws of Theorem 2.9 in [13]. Thus for any valuation w over $\{x, y\}$,

$$w(x * (x+y)) = w(x) * w(x+y) = w(x) * (w(x) + w(y)) = w(x)$$

from which we can derive (i), and similarly

$$w(x + x * y) = w(x) + w(x) * w(y) = w(x)$$

from which we can derive (ii). □

A set $U \subseteq X$ is taken to be consistent if $C_n(U) \neq X$ (compare Definition 3.9(ii) and p115 in [12], and also Sections 1.2.1 and 1.2.3). Equivalently, from Proposition 3.44, Proposition 3.42 and Corollary 3.57 below, U is consistent if for any $x \in U$, it is not the case that $x' \in U$ as well.

Corollary 3.57. *Within the logic* $\mathfrak{L}_X = ((X, +, *, {}', \mathbf{0}_X, \mathbf{1}_X), C_n)$*, for any* $x \in X$*,* $x * x' \leftrightarrow \mathbf{0}_X$*.*

Proof. By Definitions 3.27 and 3.30 and Lemma 3.51 we have

$$C_n(\{x * x'\}) = C_n(\{x, x'\}) = X = C_n(\{\mathbf{0}_X\})$$

and hence $x * x' \leftrightarrow \mathbf{0}_X$. □

Theorem 3.58 captures the familiar de Morgan Laws, as applied to $\mathfrak{L}_X$.

Theorem 3.58. *For the logic* $\mathfrak{L}_X = ((X, +, *, {}', \mathbf{0}_X, \mathbf{1}_X), C_n)$*, let* $x, y \in X$ *be any two sentences. Then,*

i) $(x + y)' \approx x' * y'$

ii) $x' + y' \approx (x * y)'$

Proof. Both results follow readily from the properties of $+$, $*$ and $'$. For any $x, y \in X$,

	$\forall u \in \mathbb{V}.[u(x) \leq_T u(x+y)$ and $u(y) \leq_T u(x+y)]$	(property of $+$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u((x+y)') \leq_T u(x')$ and $u((x+y)') \leq_T u(y')]$	(property of $'$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u((x+y)') \leq_T u(x' * y')]$	(property of $*$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u(x' * y') \in C_T(\{u((x+y)')\})]$	(property of C_T)
$\Rightarrow$	$x' * y' \in C_n(\{(x+y)'\})$	(Definition 3.30)

We also have

	$\forall u \in \mathbb{V}.[u(x' * y') \leq_T u(x')$ and $u(x' * y') \leq_T u(y')]$	(property of $*$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u(x) \leq_T u((x' * y')')$ and $u(y) \leq_T u((x' * y')')]$	(property of $'$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u(x+y) \leq_T u((x' * y')')]$	(property of $+$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u(x' * y') \leq_T u((x+y)')]$	(property of $'$ in $\mathbf{T}$)
$\Rightarrow$	$\forall u \in \mathbb{V}.[u((x+y)') \in C_T(\{u(x' * y')\})]$	(property of C_T)
$\Rightarrow$	$(x+y)' \in C_n(\{x' * y'\})$	(Definition 3.30)

By Proposition 3.34, we then have that $(x + y)' \approx x' * y'$. Similarly,

$$\begin{array}{ll} & \forall u \in \mathbb{V}.[u(x * y) \leq_T u(x) \text{ and } u(x * y) \leq_T u(y)] \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x') \leq_T u((x * y)') \text{ and } u(y') \leq_T u((x * y)')] \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x' + y') \leq_T u((x * y)')] \\ \Rightarrow & \forall u \in \mathbb{V}.[u((x * y)') \in C_T(\{u(x' + y')\})] \\ \Rightarrow & (x * y)' \in C_n(\{x' + y'\}) \end{array}$$

and

$$\begin{array}{ll} & \forall u \in \mathbb{V}.[u(x') \leq_T u(x' + y') \text{ and } u(y') \leq_T u(x' + y')] \\ \Rightarrow & \forall u \in \mathbb{V}.[u((x' + y')') \leq_T u(x) \text{ and } u((x' + y')') \leq_T u(y)] \\ \Rightarrow & \forall u \in \mathbb{V}.[u((x' + y')') \leq_T u(x * y)] \\ \Rightarrow & \forall u \in \mathbb{V}.[u((x * y)') \leq_T u(x' + y')] \\ \Rightarrow & \forall u \in \mathbb{V}.[u(x' + y') \in C_T(\{u((x * y)')\})] \\ \Rightarrow & x' + y' \in C_n(\{(x * y)'\}) \end{array}$$

and hence $x' + y' \approx (x * y)'$. □

It is natural to enquire how C_n might interact with the order $\leq_X$.

Proposition 3.59. *For any $x, y \in X$, if $y \in C_n(\{x\})$ then $x \leq_X y$.*

Proof. For any $x, y \in X$ we have

$$\begin{aligned} y \in C_n(\{x\}) \Rightarrow & \quad \forall u \in [X_0 \longrightarrow \mathfrak{T}_{\mathcal{D}}].[u(y) \in C_T(\{u(x)\})] \\ \Rightarrow & \quad \forall u \in W_{x,X_0} \cup W_{y,X_0}.[u(y) \in C_T(\{u(x)\})] \end{aligned}$$

For any $U \subseteq X_0$ with $L_x, L_y \subseteq U$, $u \in [U \longrightarrow \mathfrak{T}_{\mathcal{D}}]$ may be derived from $u_0 \in [X_0 \longrightarrow \mathfrak{T}_{\mathcal{D}}]$ by restriction from X_0 to U. It follows that for any $u \in W_{x,U} \cup W_{y,U}$, $u(y) \in C_T(\{u(x)\})$, and by Definition 2.31, $x \leq_X y$. □

We define the notion of a logical filter as follows (compare Definition 3.25 in Section 3.2.1, and also Definition 3.2 and Lemma 3.3 in [36]).

Definition 3.60. Within the logic $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$, a non-empty subset U of X is called a *logical filter (of X)* if

i) $x, y \in U$ implies $x * y \in U$

ii) $x \in U$, $y \in X$ and $y \in C_n(\{x\})$ imply $y \in U$

The set of all logical filters of X is denoted by $\mathfrak{F}_L(X)$ and carries the usual inclusion order.

By Proposition 3.59, within a logical filter, up-closure with respect to $\leq_X$ is therefore constrained by logical consequence. We shall refer to this restricted form of up-closure as "logical up-closure (with respect to $\leq_X$)".

Proposition 3.61. *For the logic $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$, let $\mathcal{C}_X = \{C_n(U) \mid U \subseteq X\}$ denote the closed sets of X under C_n. Then $\mathcal{C}_X = \mathfrak{F}_L(X)$.*

Proof. Let $U \subseteq X$ be a set of sentences. By **Cn2** of Definition 3.49, for any $x, y \in C_n(U)$, $x * y \in C_n(U)$ as well. From property **CL1** of Definition 3.3, for any $x \in C_n(U)$, $C_n(\{x\}) \subseteq C_n(U)$. From Definition 3.60, $C_n(U)$ is a logical filter and it follows that $\mathcal{C}_X \subseteq \mathfrak{F}_L(X)$. Now let U be a logical filter, and assume that $x \in C_n(U)$. Since C_n is algebraic, from Definition 3.23 we have

$$C_n(U) = \bigcup\{C_n(V) \mid V \subseteq U \text{ and } V \text{ is finite}\}$$

and so there is finite $V \subseteq U$ with $x \in C_n(V)$. By Definition 3.60, $\prod V \in U$ as well, and so $C_n(\{\prod V\}) \subseteq U$. By **Cn1** of Definition 3.49, we have $V \subseteq C_n(\{\prod V\})$, and from **CL2** and **CL3** of Definition 3.3, it follows that

$$C_n(V) \subseteq C_n(C_n(\{\prod V\})) = C_n(\{\prod V\}) \subseteq U$$

and hence $x \in U$ also, so that $C_n(U) \subseteq U$. By **CL1** of Definition 3.3, $U \subseteq C_n(U)$, so that $U = C_n(U)$ and hence $U \in \mathcal{C}_X$. Consequently, $\mathfrak{F}_L(X) \subseteq \mathcal{C}_X$, from which we obtain $\mathcal{C}_X = \mathfrak{F}_L(X)$, as required. □

Every closed set within $\mathfrak{L}_X$ therefore exhibits logical up-closure with respect to $\leq_X$.

Finally, let $X_0' = X_0 \setminus \{\mathbf{0}_X, \mathbf{1}_X\}$. With $\mathfrak{L}_X = ((X, +, *, ', \mathbf{0}_X, \mathbf{1}_X), C_n)$ in place, we impose logical independence on X_0 by requiring that

$$\forall x \in X_0'. \neg\exists U \subseteq X_0' \setminus \{x\}. [U \neq \emptyset \text{ and } (x \in C_n(U) \text{ or } x' \in C_n(U))]$$

The set X_0' is then consistent.

From Definition 2.28, $(X, +, *, ', \leq_X)$ is an ordered r-algebra of type $\mathcal{F}$ because $\leq_X$ is a quasi-order. Lemma 3.37 then establishes that the elements of X are bounded above and below by $\mathbf{1}_X$ and $\mathbf{0}_X$. Taken together, by Definition 2.28 $\mathbf{X} = (X, +, *, ', \mathbf{0}_X, \mathbf{1}_X, \leq_X)$ is then a bounded r-algebra of type $\mathcal{F}$. We shall take $\mathbf{X}$, together with the logic $\mathfrak{L}_X$ as our model of a subject. As indicated in Section 1.3, subjects play a central role in the formulation of of a B-structure, and so are an important component of the work we shall undertake in Chapter 4.

Definition 3.62. Let $\mathbf{D} = (D, \mathcal{D}_0, \mathbf{T})$ be a disposition, where $\mathbf{T} = (\mathfrak{T}_\mathcal{D}, +, *, ', \mathbf{0}, \mathbf{1}, \leq_T)$.

1. A subject of type $\mathcal{F}$ is a pair $\mathbf{S} = ((X, H, \leq_X), \mathfrak{L}_X)$ in which

   i) $(X, H, \leq_X)$ is an ordered r-algebraic space of type $\mathcal{F}$ over $\mathbf{T}$,

   ii) $\leq_X$ is derived by Definition 2.31,

   iii) H is a sheaf of $\mathfrak{T}_\mathcal{D}$-valuations over X,

   iv) $(X, +, *, ')$ is the algebra of sentences of type $\mathcal{F}$ over a non-empty, countable set X_0 of primitives,

   v) $\mathbf{X} = (X, +, *, ', \mathbf{0}_X, \mathbf{1}_X, \leq_X)$ is a bounded r-algebra of type $\mathcal{F}$,

   vi) $\mathfrak{L}_X = (\mathbf{X}, C_n)$ is a logic of type $\mathcal{F}$, where $*$ is a normal conjunction on X, $+$ is a normal disjunction on X, $'$ is a negation on X and $\mathbf{0}_X$ and $\mathbf{1}_X$ are respectively contradiction and theoremhood in $\mathfrak{L}_X$, and

   vii) C_T is the algebraic closure operator on $\mathfrak{T}_\mathcal{D}$ derived from Definition 3.27, and C_n is the algebraic closure operator on X derived from C_T through Definition 3.30.

   We refer to X as the *content* of the subject $\mathbf{S}$.

2. A morphism $((X, H_X, \leq_X), \mathfrak{L}_X) \longrightarrow ((Y, H_Y, \leq_Y), \mathfrak{L}_Y)$ of subjects is given by a morphism of ordered r-algebraic spaces from $(X, H_X, \leq_X)$ to $(Y, H_Y, \leq_Y)$ together with a weak homomorphism from $\mathfrak{L}_X$ to $\mathfrak{L}_Y$.

With regard to $\mathfrak{L}_X$, any set of sentences in X we will call a theory, and in particular we will refer to the closed sets of X as X-theories. Given $U \subseteq X$, the set $C_n(U)$ may then be thought of as the set of "dispositional consequences" of U, and under C_n the closed sets of $\mathbf{X}$ coincide with the theories of the underlying logic on $\mathbf{X}$. A theory U is consistent if $C_n(U) \neq X$. If $V \subseteq U$, we will sometimes refer to V as a sub-theory of U.

Summary

In this chapter we began to specialise the framework that we constructed in Chapter 2 by developing a model of a subject as a family of sentences equipped with an order and a logic.

We began by furnishing a family of sentences in the form of a sentence algebra of a specified type over a non-empty, countable set of primitives. The type of the algebra was chosen with the logic in mind, because we intended to use the fundamental operations as connectives in the logic.

To develop the required order, we appealed to the idea of a disposition. To compare information by a given disposition, it would be necessary to assign the disposition to the given information in the form of comparable "degrees". Had we done that directly, we would have encountered two difficulties. First, our model would have broken down as soon as the agent was discovered to hold a degree of disposition different from what we had assigned. Second, we would have required knowledge of the type used to represent the degrees of disposition. We overcame both problems by assigning these degrees of disposition non-deterministically, an idea which we drew from [8].

We therefore assigned degrees of disposition to information based on properties that a degree should satisfy, rather than based on its value. By using sets of properties or specifications rather than values, we could also retain the qualitative character of the problem. In this regard, our approach differed from models such as the transferable belief model of [46], which attempts to quantify the degree of belief of an agent explicitly, or the *ceteris paribus* networks of [7] where, although no quantification is made explicit, the qualitative preference orderings used to order the beliefs of an agent appeal directly to the values in an ordered set of preferences.

An important feature proposed in Section 1.3 was that the agent should be able to hold an "indeterminate" disposition towards a given piece of information, since it may know nothing about what that information might mean. To implement this feature, we drew on the work of several authors. In particular, from [22] and [23], we used the idea of an order derived from a set of observations about a given set of values. From [15], we drew on the idea of a bi-partite truth value to develop the surrogate degrees of disposition we would ultimately assign to each piece of information.

In the derivation of these surrogates, our intention was to allow the two components of a given surrogate to overlap – the greater the overlap, the greater the indeterminacy represented by the surrogate. Consequently, we used an idea described in [9] to develop P- and Q-specifications of degrees of disposition. The P-specification described degrees of disposition that could be held towards an item of information, while the Q-specification described degrees of disposition that could be thought of as being converse to those described by P. The interplay between these sets of degrees of disposition let us capture an element of ambivalence on the part of the agent towards the selected item of information.

For technical reasons, we ordered specifications by inclusion (rather than using a power order). More specifically, we chose the truth order of [15] because it reflected what we were after – as we moved up the order on surrogates, the P-specification would strengthen while the Q-specification weakened. Thus, the contrast between the pro- and contra-disposition represented by the surrogate would increase, strengthening the characterisation of the prescribed degrees of disposition. In contrast, under the knowledge order, the two degrees $(\emptyset, \Theta_{\mathcal{D}})$ and $(\Theta_{\mathcal{D}}, \emptyset)$ would no longer represent surrogates in which the disposition was held minimally (resp. maximally). Rather, the top and bottom elements would be $(\Theta_{\mathcal{D}}, \Theta_{\mathcal{D}})$ and $(\emptyset, \emptyset)$, which we needed to represent an indeterminate disposition.

From the order on surrogates, we could derive an order on sentences as described in Definition 2.31. To develop a logic, we appealed to the idea of "preservation of degrees of truth" described in [16]. From this idea we were able to develop a consequence function that preserved the degrees of disposition assigned to the members of a set of sentences, as set out in Definitions 3.27 and 3.30. Furthermore, the notion of consequence also had a well-defined relationship to the order $\leq_X$ on X, as shown by Propositions 3.59 and 3.61. We were then able to use the operations of our algebra as logical connectives, in the process formulating notions of conjunction, disjunction and negation.

Finally, from [36], an important requirement in the development of their belief revision relation was that, given a sentence x, any x-compatible set of sentences should be contained in a maximal x-compatible set of sentences. We were able to meet this requirement by ensuring that our consequence function was algebraic.

Chapter 4

Belief Revision by Disposition

In Section 1.3, we proposed to replace the epistemic entrenchments used in [36] with a more general quasi-order. These epistemic entrenchments played a central role in the development of the belief revision relations of [36]. Our motivation for replacing them with a quasi-order was to extend the model of [36] to obtain a more general version of relational belief revision where sentences are ordered by how strongly they exhibit a given property. Accordingly, theories are then revised by how strongly the sentences exhibit the property.

We developed the quasi-order from a propositional attitude or disposition in Chapter 3, and based on a given disposition we were also able to construct a logic on a family of sentences. We now continue this development by formally exhibiting a belief revision relation that is derived from a disposition, and hence uses a more general quasi-order. We then describe how a B-structure may be constructed from the components we have developed in Chapters 2 and 3, and conclude by providing a worked example to show how the theory we have developed may be applied to related areas of research. We proceed as follows.

Chapter Guide:

Section 4.1: Belief Revision. In this section, we develop a belief revision relation for a family of sentences that is ordered by a disposition. The proposed relation is formulated in terms of belief bases rather than belief sets. We then show how a B-structure may be constructed from the components set out in Chapters 2 and 3. This B-structure is relational in that the underlying subjects are equipped with a relational belief revision operation. We conclude with a brief discussion of how belief revision may be conducted with a relational B-structure.

Section 4.2: Towards an Application of Our Work. In [44], the author presents a "phenomenal, dispositional account of belief". This model provides an innovative approach to belief in which the beliefs held by an agent are portrayed in terms of what are called dispositional stereotypes. The approach successfully accounts for those cases where an agent can neither be described as believing or not believing something. In this section, we describe this approach, focussing in particular on dispositional stereotypes. We then show how the theory we have developed for our work may be applied to the approach.

4.1 Belief Revision

We would now like to develop a belief revision relation by using the dispositions that we derived in Chapter 3. Our exposition in this sub-section mostly follows that of [36], and through it we are able to demonstrate how the theory that we have developed can be applied to the work of [36].

4.1.1 A Belief Revision Relation in the Style of Lindström and Rabinowicz

Let $\mathbf{S} = ((X, H, \leq_X), \mathfrak{L}_X)$ be a subject of type $\mathcal{F} = \{+, *, '\} = \{2, 2, 1\}$, for which $\leq_X$ is derived from a disposition $\mathbf{D} = (D, \mathcal{D}_0, \mathbf{T})$ and $\mathfrak{L}_X = (\mathbf{X}, C_n)$, where C_n is given by Definition 3.30. For C_n, let $\mathcal{C}_X$ denote the family of closed sets of X. We take the "beliefs" of an agent to be represented by the set $G \subseteq X$ of sentences. In keeping with our convention, G is considered a theory. We shall assume that this theory is consistent so that $C_n(G) \neq X$, that G is not necessarily an X-theory, and that the surrogate degree of disposition towards each member of G can be determined with the reference assignment v of the agent.

For belief revision, we shall restrict our consideration to sentences that are finite in the following sense.

Definition 4.1. Let $\mathcal{F}$ be a language of algebras, and let $\mathbf{X} = (X, F)$ be the sentence algebra of type $\mathcal{F}$ over a non-empty, countable set X_0 of primitives. For any $x \in X$, we say that x is a *finite sentence* if L_x is finite, and for any $y \in L_x$, y appears in x only finitely many times.

Given a new finite sentence y, the agent will seek to revise G to include y. The outcome of this revision depends on whether the agent can accommodate y amongst the sentences in G, given the logic $\mathfrak{L}_X$ and its reference commitment v. If it cannot, we expect that it will abandon some of its existing sentences to accommodate the new information, effectively reverting to a sub-theory of G. We would like this retiring of sentences to be done so that the least number of sentences is given up to accommodate y. Intuitively, we seek a sub-theory $K \subseteq G$ that is maximally compatible with y, as described in Definition 3.45. Fortunately, by Proposition 3.47, such a sub-theory of G does exist.

We propose a belief revision relation as follows.

Definition 4.2. For any $G, H \in \mathcal{C}_X$ let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$. For any $y \in X$, let $R_y \subseteq \mathcal{C}_X \times \mathcal{C}_X$ be a binary relation on $\mathcal{C}_X$ such that for any $G, H \in \mathcal{C}_X$, GR_yH if

i) $\vdash_{\mathfrak{L}_X} y'$ and $H_0 = \{\mathbf{0}_X\}$, or

ii) G_0 contains a possibly empty, maximal y-compatible sub-theory K and $H_0 = K \cup \{y\}$

If GR_yH, we say that H *is a revision of* G *with* y *relative to* v.

This definition may be compared to Definition 4.1 in [36]. The relation R_y between the belief sets G and H is determined by properties of the underlying belief bases G_0 and H_0, and we may think of R_y as inducing a binary relation R^0_y on $\mathcal{P}(X)$, where $G_0 R^0_y H_0$ if and only if GR_yH. We may then regard R_y as an implicit and R^0_y as an explicit belief revision relation (compare Remark 1.9).

We claim that the relation R_y of Definition 4.2 is in fact a belief revision relation. That is, R_y satisfies the seven LR postulates **R1–R7** listed in Section 1.2.3. For easy reference, we restate these axioms here, reformulated to use the notation we have established thus far. In the list set out below, G and H are X-theories and $y, z \in X$.

R1: There exists $H \in \mathcal{C}_X$ such that $H \in R_y(G)$

R2: If $H \in R_y(G)$, then $y \in H$

R3: If $y' \notin G$ and $H \in R_y(G)$, then $H = C_n(G \cup \{y\})$

R4: If $\nvdash_{\mathfrak{L}_X} y'$ and $H \in R_y(G)$, then $\mathbf{0}_X \notin H$

R5: If $\vdash_{\mathfrak{L}_X} y \leftrightarrow z$, then $H \in R_y(G)$ if and only if $H \in R_z(G)$

R6: If $H \in R_y(G)$ and $z' \notin H$, then $C_n(H \cup \{z\}) \in R_{y*z}(G)$

R7: If $H \in R_y(G)$ and for all $L \in \mathcal{C}_X$ we have that if $L \in R_{y+z}(G)$ then $y' \notin L$, then there exists $L \in \mathcal{C}_X$ such that $L \in R_{y+z}(G)$ and H is given by $C_n(L \cup \{y\})$

It follows readily from Definition 4.2 that R_y satisfies Axioms **R1** and **R2**, so we state these two claims without proof.

Lemma 4.3. *The relation R_y satisfies Axiom* **R1***.*

Lemma 4.4. *The relation R_y satisfies Axiom* **R2***.*

Lemma 4.5. *The relation R_y satisfies Axiom* **R3***.*

Proof. Let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$, and assume that $y' \notin G$ and $H \in R_y(G)$. Since $y' \notin G$, we do not have $\vdash_{\mathfrak{L}_X} y'$ and hence Definition 4.2(i) does not apply. Furthermore, because $y' \notin G$ it follows that $y' \notin G_0$ and by Definition 3.45 G_0 is a maximal y-compatible sub-theory of itself, so that $H_0 = G_0 \cup \{y\}$. Applying Lemma 3.50, we then obtain

$$H = C_n(H_0) = C_n(G_0 \cup \{y\}) = C_n(C_n(G_0) \cup \{y\}) = C_n(G \cup \{y\})$$

as required. □

Lemma 4.6. *The relation R_y satisfies Axiom* **R4***.*

Proof. Let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$, and assume that $\nvdash_{\mathfrak{L}_X} y'$ and $H \in R_y(G)$. Because $\nvdash_{\mathfrak{L}_X} y'$, Definition 4.2(i) does not apply and we may write $H_0 = K \cup \{y\}$ where K is a maximal y-compatible sub-theory of G_0. From Proposition 3.44, because $y' \notin C_n(K)$, $\mathbf{0}_X \notin K$. Hence $\mathbf{0}_X \notin K \cup \{y\}$, and by Proposition 3.44 we have $\mathbf{0}_X \notin C_n(K \cup \{y\}) = H$. □

Lemma 4.7. *The relation R_y satisfies Axiom* **R5***.*

Proof. Let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$, and assume that $\vdash_{\mathfrak{L}_X} y \leftrightarrow z$. If $H_{0,y} = \{\mathbf{0}_X\}$, then we must have $\vdash_{\mathfrak{L}_X} y'$. By Definition 3.40 and Lemma 3.43, if $\vdash_{\mathfrak{L}_X} y \leftrightarrow z$ then $\vdash_{\mathfrak{L}_X} y' \leftrightarrow z'$, and from $\vdash_{\mathfrak{L}_X} y'$ we therefore have $\vdash_{\mathfrak{L}_X} z'$. But then $H_{0,z} = \{\mathbf{0}_X\}$ as well. It follows that if $H = X = C_n(H_{0,y}) \in R_y(G)$ then also $X \in R_z(G)$. Otherwise, let K be a maximal y-compatible sub-theory of G_0, so that $H_{0,y} = K \cup \{y\}$. Then

$$\begin{aligned} z' \in C_n(K) \Rightarrow\ & C_n(\{z'\}) \subseteq C_n(K) && \text{(property CL2 of Definition 3.3)} \\ \Rightarrow\ & C_n(\{y'\}) \subseteq C_n(K) && \text{(since } \vdash_{\mathfrak{L}_X} y \leftrightarrow z\text{)} \\ \Rightarrow\ & y' \in C_n(K) && \text{(property CL1 of Definition 3.3)} \end{aligned}$$

and hence K could not have been a y-compatible sub-theory of G_0, and we reach a contradiction. Thus, K is z-compatible. Next suppose that there is $K' \subseteq G_0$ such that K' is z-compatible and $K \subset K'$. Then

$$\begin{aligned} K \subset K' \text{ and } z' \notin C_n(K') \Rightarrow\ & C_n(\{z'\}) \not\subseteq C_n(K') \\ \Rightarrow\ & C_n(\{y'\}) \not\subseteq C_n(K') && \text{(since } \vdash_{\mathfrak{L}_X} y \leftrightarrow z) \\ \Rightarrow\ & y' \notin C_n(K') \end{aligned}$$

and hence K could not have been a maximal y-compatible set and again we reach a contradiction. Thus K is a maximal z-compatible sub-theory of G_0, and $H_{0,z} = K \cup \{z\}$. Applying Lemma 3.50, we then obtain

$$C_n(K \cup \{y\}) = C_n(K \cup C_n(\{y\})) = C_n(K \cup C_n(\{z\})) = C_n(K \cup \{z\})$$

and it follows that if $H = C_n(K \cup \{y\}) \in R_y(G)$ then also $H \in R_z(G)$. By similar analysis, we can show that if $H \in R_z(G)$ then also $H \in R_y(G)$, which then establishes the result. □

Lemma 4.8. *The relation R_y satisfies Axiom* **R6**.

Proof. Let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$, and assume that GR_yH and $z' \notin H$. Since $z' \notin H$, we have $H \neq X$ and hence $\nvdash_{\mathfrak{L}_X} y'$, so that Definition 4.2(i) does not apply to R_y. There is therefore a maximal y-compatible sub-theory K of G_0 with $H = C_n(H_0) = C_n(K \cup \{y\})$. By hypothesis, $z' \notin H$. Suppose that K is not $(y * z)$-compatible. Then $(y * z)' \approx y' + z' \in C_n(K)$ and hence also $y' + z' \in C_n(K \cup \{y\})$, from which we have $C_n(K \cup \{y\} \cup \{y' + z'\}) = C_n(K \cup \{y\})$. Now,

$$\begin{aligned} & C_n(K \cup \{y\} \cup \{y' + z'\}) \\ =\ & C_n(K \cup \{y\} \cup \{y'\}) \cap C_n(K \cup \{y\} \cup \{z'\}) && \text{(Lemma 3.51(ii))} \\ =\ & X \cap C_n(K \cup \{y\} \cup \{z'\}) && \text{(Definitions 3.30 and 3.27)} \\ =\ & C_n(K \cup \{y\} \cup \{z'\}) \end{aligned}$$

and hence $C_n(K \cup \{y\}) = C_n(K \cup \{y\} \cup \{z'\})$ in which case $z' \in C_n(K \cup \{y\}) = H$, contrary to our hypothesis. Thus K is $(y * z)$-compatible. Next, suppose that K is not a maximal $(y * z)$-compatible sub-theory of G_0. Then there is $K' \subseteq G_0$ with $K \subset K'$ and $(y * z)' \approx y' + z' \notin C_n(K')$. By **Dn1** of Definition 3.49, $y' \notin C_n(K')$ either, and hence K could not have been a maximal y-compatible sub-theory of G, and we reach a contradiction. It follows that K is also a maximal $(y * z)$-compatible sub-theory of G_0 and hence that $GR_{y*z}H'$ with $H'_0 = K \cup \{y * z\}$. As before, by Lemma 3.51(i) and Lemma 3.50, $H' = C_n(K \cup \{y, z\}) = C_n(H \cup \{z\})$. □

Lemma 4.9. *The relation R_y satisfies Axiom* **R7**.

Proof. Let $G_0, H_0 \subseteq X$ be such that $G = C_n(G_0)$ and $H = C_n(H_0)$, and assume that GR_yH and that for any X-theory L, if $GR_{y+z}L$ then $y' \notin L$. Now,

$$\begin{aligned} & \exists L \subseteq \mathcal{C}_X.[GR_{y+z}L] && \text{(by Axiom } \mathbf{R1}) \\ \Rightarrow\ & y' \notin L && \text{(by hypothesis)} \\ \Rightarrow\ & \nvdash_{\mathfrak{L}_X} y' && \text{(by Definition 3.36)} \end{aligned}$$

and hence Definition 4.2(i) does not apply to R_y. Thus, we have GR_yH with $H = C_n(K \cup \{y\})$ and K is a maximal y-compatible sub-theory of G_0. By Condition **Dn1** of Definition 3.49, since $y \in H$, $y + z \in H$ also, and because $H \neq X$, by Propositions 3.44 and 3.57, $(y + z)' \notin H$ so that K is $(y + z)$-compatible. Suppose that K is not a maximal $(y + z)$-compatible sub-theory of G_0. Then there is $K' \subseteq G_0$ such that $K \subset K'$ and K' is a maximal $(y + z)$-compatible sub-theory of G_0. Taking $L_0 = K' \cup \{y + z\}$ and $L = C_n(L_0)$, we then have $GR_{y+z}L$. But since K was a maximal y-compatible sub-theory of G_0, $y' \in C_n(K') \subseteq L$, contrary to our hypothesis, and so K is also a maximal $(y+z)$-compatible sub-theory of G_0. Thus, taking $L_0 = K \cup \{y+z\}$ and $L = C_n(L_0)$, we then have $GR_{y+z}L$ and $H = C_n(K \cup \{y\})$. By applying Proposition 3.56, we find that

$$\begin{aligned} H = \; & C_n(K \cup \{y\}) \\ = \; & C_n(K \cup \{y * (y+z)\}) \\ = \; & C_n(K \cup \{y, y+z\}) && \text{(Lemma 3.51(i))} \\ = \; & C_n(L_0 \cup \{y\}) \\ = \; & C_n(C_n(L_0) \cup \{y\}) && \text{(Lemma 3.50)} \\ = \; & C_n(L \cup \{y\}) \end{aligned}$$

which gives us the result, as required. □

The preceding lemmas show that the relation R_y as set out in Definition 4.2 satisfies the axioms that define a belief revision relation, and hence R_y is in fact a belief revision relation.

Theorem 4.10. *The relation R_y of Definition 4.2 is a belief revision relation.*

We have now presented a scheme under which a belief revision relation can be derived for a family of sentences that has been ordered by a given disposition $\mathbf{D}$. In the next section, we are able to build a B-structure and show how relational belief revision may be carried out in it.

4.1.2 Relational B-Structures

We let $\mathfrak{S} = \{\mathbf{S}_j\}_{j \in J}$ be a countable family of subjects of type $\mathcal{F} = \{+, *, '\} = \{2, 2, 1\}$. Recalling Definition 3.62, for each $j \in J$, $\mathbf{S}_j = ((X_j, H_j, \leq_j), \mathfrak{L}_j)$ in which

i) $(X_j, H_j, \leq_j)$ is an ordered r-algebraic space of type $\mathcal{F}$ over $\mathbf{T}_j = (\mathfrak{T}_j, +, *, ', \mathbf{0}_{T_j}, \mathbf{1}_{T_j}, \leq_{T_j})$,

ii) $\leq_j$ is derived from $\mathbf{T}_j$ by Definition 2.31,

iii) H_j is a sheaf of $\mathbf{T}_j$-valuations over X_j,

iv) $(X_j, +, *, ')$ is the algebra of sentences of type $\mathcal{F}$ over a non-empty, countable set $X_{0,j}$ of primitives and X_j is the content of subject $\mathbf{S}_j$,

v) $\mathbf{X}_j = (X_j, +, *, ', \mathbf{0}_j, \mathbf{1}_j, \leq_j)$ is a bounded r-algebra of type $\mathcal{F}$,

vi) $\mathfrak{L}_j = (\mathbf{X}_j, C_j)$ is a logic of type $\mathcal{F}$, where $*$ is a normal conjunction on X_j, $+$ is a normal disjunction on X_j, $'$ is a negation on X_j and $\mathbf{0}_j$ and $\mathbf{1}_j$ are respectively contradiction and theoremhood in $\mathfrak{L}_j$, and

vii) C_j is derived via Definitions 3.27 and 3.30

Each $\mathbf{T}_j$ originates from a disposition $\mathbf{D}_j = (D_j, \mathcal{D}_{0,j}, \mathbf{T}_j)$. To allow more than one disposition to be applied to the content X_j of a subject $\mathbf{S}_j$, $\mathfrak{S}$ may contain copies of $\mathbf{S}_j$ for which the order is induced by a different disposition. Correspondingly, the same disposition may be applied to more than one member of $\mathfrak{S}$. Consequently, it is possible that for $\mathbf{S}_j, \mathbf{S}_k \in \mathfrak{S}$ we have $X_j = X_k$ but $\leq_j \neq \leq_k$, and it is also possible that although $X_j \neq X_k$, $\mathbf{D}_j$ and $\mathbf{D}_k$ are the same disposition.

To simplify our work, for distinct $i, j \in J$, we shall treat $X_{0,i}$ and $X_{0,j}$ as being disjoint. Consequently, for any $\mathbf{S}_j, \mathbf{S}_k \in \mathfrak{S}$ for which $X_j = X_k$, we nonetheless treat X_j and X_k as distinct sets of sentences. We let $X_0 = \bigcup_{j \in J} X_{0,j}$, and for each $\mathbf{S}_j$ we assume that an agent holds the beliefs in T_j, corresponding to the belief base Γ_j.

As described in Section 2.3.2, we combine the $\mathbf{T}_j$ by forming a coalesced sum of the $\mathfrak{T}_j$, identifying the top and bottom elements in the process, and in this way we obtain a new bounded r-algebra $\mathbf{T} = (\mathfrak{T}, +, *, ', \mathbf{0}_T, \mathbf{1}_T, \leq_T)$. For each $j \in J$, H_j is now a sheaf of $\mathbf{T}$-simulations. Taking U_0 to be a countable set, for each $j \in J$, we let $U_j \subseteq U_0$ and $\phi_j : U_j \longrightarrow X_j$ be such that ϕ_j maps U_j bijectively to X_j, and we let $\mathcal{U} = \{U_j\}_{j \in J}$ and $X = \bigcup_{j \in J} U_j$. Taking $\mathsf{M} = \{(X_j, H_j, \leq_j)\}_{j \in J}$, we construct an ordered manifold $\mathbf{X} = (X, H_X, \leq_X)$ of type M as set out in Section 2.3.2. Borrowing notation from Section 2.3.2, for each $j \in J$, if $x \in U_j$ we take $\phi_j(x)$ to be a sentence that expresses or represents x in $\mathbf{S}_j$, and we let $V_j = \phi_j^{-1}(\Gamma_j)$ and $V_j^+ = \phi_j^{-1}(T_j)$.

By combining each U_j with the corresponding V_j^+, we may then form $\mathcal{B} = \{(U_j, V_j^+)\}_{j \in J}$ as a B-structure on X_0 (compare Definition 1.8). For any $j \in J$, the subject $\mathbf{S}_j$ is accessible via the corresponding coordinate chart (U_j, ϕ_j), and in this sense $\mathbf{S}_j$ participates in $\mathcal{B}$. Consequently, we name $\mathcal{B}$ a relational B-structure (of type M) because each subject that participates in $\mathcal{B}$ is equipped with a relational belief revision operation.

Let y be a finite sentence of type $\mathcal{F}$, formed from the primitives in X_0, and let L_y be the smallest language of y (compare Lemma 1.4, Definition 1.5 and Remark 1.6). To revise $\mathcal{B}$ with y, we shall apply the "Option B Revision" procedure (compare Section 1.2.2) as follows:

1. *We form a new subject by combining those subjects that are affected by the new information y.*

   Let $J_y = \{j \in J \mid L_y \cap X_{0,j} \neq \emptyset\}$, and let $X_{0,y} = \bigcup_{j \in J_y} X_{0,j}$. Let $(X_y, +, *, ')$ be the algebra of sentences of type $\mathcal{F}$ over $X_{0,y}$. Via Definition 2.31 we can derive an order $\leq_y$ on X_y to form the bounded r-algebra $\mathbf{X}_y = (X_y, +, *, ', \leq_y)$ of type $\mathcal{F}$. The Alexandroff topology ΩX_y on X and the sheaf H_y of $\mathbf{T}$-simulations over X_y may then be derived as set out in Section 2.2.2 to give us the ordered space $(X_y, H_y, \leq_y)$. From the work of Section 3.2 we can construct a logic $\mathfrak{L}_{X_y} = (\mathbf{X}_y, C_y)$, where C_y is developed as described in Definitions 3.27 and 3.30. Taken together, we then have the subject $\mathbf{S}_y = ((X_y, H_y, \leq_y), \mathfrak{L}_{X_y})$.

2. *We determine the beliefs about the subject $\mathbf{S}_y$.*

   For each $j \in J_y$, $X_{0,j} \subseteq X_{0,y}$. Consequently, for any $j \in J_y$ the algebra $(X_j, +, *, ')$ can be embedded in $(X_y, +, *, ')$. Letting $f_j : X_j \longrightarrow X_y$ represent the embedding, the beliefs about $\mathbf{S}_y$ are given by $T'_y = C_y(\Gamma'_y)$, where $\Gamma'_y = \bigcup_{j \in J_y} f_j(\Gamma_j)$. We then revise T'_y with the new information y as described in Section 4.1.1 (compare Definition 4.1.1) to obtain the new belief base Γ_y and beliefs T_y.

3. *We reconstitute the manifold $\mathbf{X}$.*

   We remove from M all of the ordered spaces that correspond to subjects affected by y, and we add the ordered space of the new subject $\mathbf{S}_y$ to form the class

   $$\mathsf{M}' = \mathsf{M} \setminus \{(X_j, H_j, \leq_j)\}_{j \in J_y} \cup \{(X_y, H_y, \leq_y)\}$$

   of ordered spaces. We then form a new manifold $\mathbf{X}'$ of type M' and a new relational B-structure $\mathcal{B}'$ as before.

If instead we wished only to know whether y followed from the beliefs in $\mathcal{B}$, we follow the same procedure to construct the new subject $\mathbf{S}_y$ and the corresponding belief base Γ'_y (steps 1 and 2 above). Following Definition 1.12, $v_{\mathcal{B}} : X_y \longrightarrow \{\bot, \mathsf{true}, \mathsf{false}, \top\}$ is such that if Γ'_y is consistent, then

i) if $\Gamma'_y \vdash_{\mathfrak{L}_y} y$, then $v_{\mathcal{B}}(y) = \mathsf{true}$,

ii) if $\Gamma'_y \vdash_{\mathfrak{L}_y} y'$, then $v_{\mathcal{B}}(y) = \mathsf{false}$, and

iii) $v_{\mathcal{B}}(y) = \perp$ otherwise

and if Γ'_y is inconsistent, then $v_{\mathcal{B}}(y) = \top$.

4.2 Towards an Application of Our Work

In [44], the author presents a "phenomenal, dispositional account of belief" in which the beliefs held by an agent are characterised by what are called dispositional stereotypes. An interesting feature of the approach is that, without having to use a measure of uncertainty such as a probability, it can successfully account for cases where an agent can neither be described as believing nor disbelieving something. In this section, we describe this approach, focussing in particular on these dispositional stereotypes. To demonstrate the applicability of our framework to other research, we then show how the tools we have developed may be applied to this aspect of the approach.

4.2.1 A Dispositional Account of Belief

As described in Section 1.1, we may see belief as referring to the representation held by an agent of the truth value of a proposition [24]. These representations have the character of discrete entities, either fully present or wholly absent in the "mind" of the agent [44]. This view, however, disagrees with the view that belief comes in degrees of uncertainty, which in turn does not handle cases where an agent lies somewhere between believing and disbelieving a proposition. To overcome these difficulties, [44] provides a phenomenal, dispositional account of belief in which conscious experience plays a central role and belief is treated as being disposed to experience and do certain things.

The key element of the account is the notion of dispositional stereotype. Stereotypes are sets of properties normally associated with an object. The accuracy of a stereotype depends on how well it describes the objects to which it is normally applied. Intuitively, certain properties may be associated with an object more frequently than others, and we then have the notion that these properties are more central to a description of the object while others are more peripheral. These central properties allow us to agree, for example, on what a typical object of a given type would be.

A dispositional stereotype is then just a stereotype that consists of what are called dispositional properties. As described in [44], a dispositional property is a conditional statement whereby an object enters a state in which the property is manifested if a particular trigger condition is met. To illustrate, the dispositional stereotype of being hot-tempered might include dispositional properties such as responding angrily (the manifestation) to minor provocations (the trigger condition), quickly expressing frustration or responding with aggression when one's will is thwarted.

The approach of [44] is then to assign a dispositional stereotype to belief in a given proposition. An agent that presents all of the properties in the stereotype for believing a proposition ϕ can then be taken to believe that ϕ is the case, while an agent that presents none of these properties can not. Informally, the more properties in the stereotype that are presented by the agent, and the more central these properties are to the stereotype, the more appropriate it is to describe the agent as believing that ϕ. In this way, the approach can account for a variety of "in-between" cases where an agent neither believes nor disbelieves that ϕ.

To substantiate, the author describes two case studies where an agent m appears to believe that ϕ in certain situations, and that $\neg\phi$ in other situations. It is then not clear what m believes about ϕ. In the author's view, degrees of belief do not adequately describe these "in-between" cases because it is not that m has little confidence that ϕ is true, but rather that in some situations m is confident that ϕ is true, whereas in other situations m is confident that ϕ is false (compare p13 in [44]). The advantage of the approach is that, from the stereotype presented in a given situation, a more definite answer as to what m believes about ϕ can be given.

The dispositions of [44] are different from the propositional attitudes we described in Chapter 1 and sought to model in Chapter 3. We referred to these attitudes as dispositions in keeping with the sense of the term "propositional attitude" as used in [45]. There are, however, some similarities between our work and that of [44], and to avoid confusion we shall refer to the dispositional stereotypes of [44] as stereotypes, reserving the term "disposition" for the structure set out in Definition 3.20. Over the next two sub-sections, we shall examine how the stereotypes of [44] may be modelled by the dispositions of Chapter 3.

4.2.2 Dispositions and Stereotypes

We start with a stereotype S_0, represented by a finite set of properties. We take S_0 to represent a characteristic β, so that β is true of an entity x if x presents all of the properties in S_0. Usually, x will not present all of these properties, so it is only appropriate to some extent to say that β is true of x.

We would like to characterise this extent, and to this end, let D be a finite sample population whose members exhibit the properties in S_0. We take S_0 and D to be such that for any $\theta \in S_0$, there is $d \in D$ such that d presents θ, for any $d \in D$ there is $\theta \in S_0$ such that d presents θ, and for any two distinct members of D there is a property in S_0 that is presented by one and not the other. We also assume a property $\theta \in S_0$ such that every $d \in D$ presents θ. For each $\theta \in S_0$ we have $U_\theta = \{d \in D \mid d \text{ presents } \theta\}$, and the family $\mathcal{D}_0 = \{U_\theta\}_{\theta \in S_0}$ of subsets of D is then an admissible family of observations about D (Definition 3.1).

Given $X \subseteq D$, we may ask which property of S_0 best describes the members of X. We could approximate X "from below" and form a new property as $X' = \bigcup\{U \in \mathcal{D}_0 \mid U \subseteq X\}$, but this could leave some members of X unaccounted for and instead we take $X' = \bigcap\{U \in \mathcal{D}_0 \mid X \subseteq U\}$. To ensure that X' is itself a property of S_0, we form the topped $\bigcap$-structure $\mathcal{D}$ on D by closing $\mathcal{D}_0$ under non-empty intersections (compare the discussion following Definition 3.1). We write $\Theta_{\mathcal{D}}$ for the corresponding set of properties and take this set to represent our supplemented stereotype which we now denote by S.

Given $h : D \longrightarrow \mathcal{P}(\mathcal{D}) : d \mapsto \{U \in \mathcal{D} \mid d \in U\}$ we may order the members of D as

$$c \preceq d \text{ if and only if } h(c) \subseteq h(d)$$

for any $c, d \in D$, as with Definition 3.6. This order just means that d is a better match for S than c since it presents more properties of S than c does. By Proposition 3.9, d then satisfies a stronger property than c and relative to S, d improves c. Equivalently, if $c \preceq d$ then S describes d more accurately than it does c, and in view of [44] it is therefore more appropriate to say that β is true of d than that it is true of c. As in Section 3.1.1, we may artificially include top and bottom elements $d_{\mathbf{1}}$ and $d_{\mathbf{0}}$ in D if these are not already present. In particular, the element $d_{\mathbf{1}}$ exhibits all of the properties in S, so it is fully appropriate to say that β is true of $d_{\mathbf{1}}$.

For any $\phi, \psi \in \Theta_{\mathcal{D}}$, the inclusion order on the properties in $\Theta_{\mathcal{D}}$ gives us

$$\phi \sqsubseteq_i \psi \text{ if and only if } U_\phi \subseteq U_\psi$$

Under this order, properties that are higher up are more central to the stereotype, since they describe more members of D. Similarly, for any $\phi, \psi \in \Theta_{\mathcal{D}}$ the refinement order on $\Theta_{\mathcal{D}}$ gives us

$$\phi \sqsubseteq_r \psi \text{ if and only if } U_\phi \preceq^+ U_\psi$$

We regard $\sqsubseteq_r$ as a form of refinement, since if $\phi \sqsubseteq_r \psi$ then any $c \in U_\phi$ is improved by at least one $d \in U_\psi$ and every $d \in U_\psi$ improves at least one $c \in U_\phi$. Given our interpretation of $\preceq$, we may also think of ψ as a better characterisation of S than ϕ, and in this way, the order $\sqsubseteq_r$ provides us with additional insight into the nature of S. As in Section 3.1.1, to $\Theta_{\mathcal{D}}$ we add the property $\theta_{\mathbf{0}} = \{d_{\mathbf{0}}\}$ ($\theta_{\mathbf{1}} = \{d_{\mathbf{1}}\}$ is already present).

A subset Θ of $\Theta_{\mathcal{D}}$ determines a partial stereotype of S. In particular, under $\sqsubseteq_r$, the partial stereotype $\uparrow\theta$ represents a set of properties, all of which characterise S at least as well as θ. Correspondingly, $(\uparrow\theta)^c$ contains those properties, none of which characterise S as well as θ. As in Section 3.1.1, we may then form a family of partial stereotypes by closing $\mathcal{S}_0 = \{\uparrow\theta \mid \theta \in \Theta_{\mathcal{D}}\} \cup \{(\uparrow\theta)^c \mid \theta \in \Theta_{\mathcal{D}}\}$ under non-empty intersections and arbitrary unions to form a topped $\bigcap$-structure on $\Theta_{\mathcal{D}}$. We may then form the set $\mathfrak{T}_{\mathcal{D}}$ of $\mathcal{D}$-surrogates as set out in Definition 3.13.

Suppose that we now encounter an entity $x \notin D$ that presents the properties in the partial stereotype $\Theta \subseteq \Theta_{\mathcal{D}}$. This partial stereotype gives rise to a $\mathcal{D}$-surrogate $t = (P_\Theta, Q_\Theta) \in \mathfrak{T}_{\mathcal{D}}$. Based on the discussion preceding Definition 3.13, P_Θ is just the smallest closed stereotype to contain Θ, whereas Q_Θ may be considered as a closed counter-stereotype because it was formed by considering those properties not presented by x. As before, in general it is possible to find members of D that satisfy both P_Θ and Q_Θ, so (P_Θ, Q_Θ) can embody an element of ambivalence with regard to S.

We may take the position of t in the order on $\mathfrak{T}_{\mathcal{D}}$ as an indication of how appropriate it is to say that β is true of x. Depending on the situation, x may present a different set of properties. This gives rise to a family $\{\Theta_i\}_{i \in I}$ of subsets of $\Theta_{\mathcal{D}}$, where in situation i, x presents the properties in Θ_i. Each partial stereotype Θ_i gives rise to a $\mathcal{D}$-surrogate t_i, and in turn we may take t_i as an indication of how appropriate it is to say that β is true of x in situation i.

Finally, suppose that we have a stereotype $\overline{S}$ that represents the characteristic $\neg\beta$, and that in situation i, x presents $\overline{S}$ to the extent represented by the $\mathcal{D}$-surrogate $\bar{t}_i$. We can then form the pair $(t_i, \bar{t}_i)$ to represent the extent to which β and $\neg\beta$ are true of x in situation i. If we apply the truth order of Section 3.1.2 to these pairs, we then have

$$(t, \bar{t}) \leq (t', \bar{t}') \quad \text{if} \quad t \leq t' \text{ and } \bar{t}' \leq \bar{t}$$

The pair $(\mathbf{0}, \mathbf{1})$ indicates that $\neg\beta$ is wholly true of x, and correspondingly $(\mathbf{1}, \mathbf{0})$ indicates that β is wholly true of x. The remaining pairs correspond to "in-between" cases where neither β nor $\neg\beta$ is completely true of x.

We can add more subtlety to our approach by viewing the properties in S as a (propositional) language from which other complex stereotypical properties can be constructed. The set of properties actually presented by an entity x is then just a theory of this language. The properties in the theory may be presented by x at most to some degree, and based on the work of Section 2.2.2 we can then explore to what extent this theory is representative of the stereotype and also to what extent β is true of x. We may also consider what happens when x presents a new stereotypical property. In the next section, we explore these aspects in more detail.

4.2.3 Dispositional Order, Revision and Stereotypes

We begin once more with a stereotype S_0, represented by a finite set of properties. We take S_0 to represent a characteristic β, so that β is true of an entity x if x presents all of the properties in S_0. Let

$\mathcal{F} = \{+, *, '\} = \{2, 2, 1\}$ be a language of algebras, and let $\mathbf{S} = (S, +, *, ')$ be the sentence algebra of type $\mathcal{F}$ over S_0. The carrier set S of $\mathbf{S}$ contains more complex stereotypical properties that can be attributed to β, and subsets of S represent partial stereotypes that an entity could present. Given an entity that presents the properties in $X \subseteq S$, we would once more like to determine the extent to which β is true of the entity.

To this end, suppose we are given a bounded r-algebra $\mathbf{T} = (\mathfrak{T}, +, *, ', \mathbf{0}_T, \mathbf{1}_T, \leq_T)$ of type $\mathcal{F}$. We would like the values in $\mathfrak{T}$ to indicate the extent to which a property $s \in S$ is exhibited by an entity x. To this end, a T-assignment over S_0 will allow us to compute a value in $\mathfrak{T}$ for s. Suppose then that x exhibits the properties in S_0 to the extents represented by the reference assignment v. From Proposition 2.29, the T-assignments form the bounded r-algebra $([S_0 \longrightarrow \mathfrak{T}], +, *, ', v_\mathbf{0}, v_\mathbf{1}, \leq_{S_0})$. In particular, $v_\mathbf{1}$ assigns $\mathbf{1}_T$ to any member of S_0, so if $v = v_\mathbf{1}$ then β is completely true of x. Correspondingly, if $v = v_\mathbf{0}$ then x does not exhibit β at all. For any other reference assignment, β is only partly true of x.

We may think of $v_\mathbf{1}$ as an idealised state in which β is completely true of x (the corresponding element was called "THE TRUTH" in Chapter 4 of [8]). Consequently, as we move up along $\leq_{S_0}$, x approaches this idealised state, and β is increasingly true of x. The members of $[S_0 \longrightarrow \mathfrak{T}]$ become more like $v_\mathbf{1}$, which gives us a notion of "similitude" that colloquially we may think of as "S_0-likeness". In this sense, assignments provide an indication of the extent to which x presents S_0 and hence to which β is true of x.

We can use this idea to develop an order on the members of S. From Section 2.2.2, recall that in the setting of valuations and assignments, the sets $W_{s,U}$ and $W_{X,U}$ of models of $s \in S$ and $X \subseteq S$ given the context U have the form

$$\begin{aligned} W_{s,U} &= \{w \in [U \longrightarrow \mathfrak{T}] \mid w \models_v s\} \quad \text{for any } U \subseteq S_0 \text{ with } L_s \subseteq U \\ W_{X,U} &= \{w \in [U \longrightarrow \mathfrak{T}] \mid w \models_v X\} \quad \text{for any } U \subseteq S_0 \text{ with } L_X \subseteq U \end{aligned}$$

where L_s and L_X are the languages of s and X (see Definitions 2.33 and 2.30).

Definition 4.11. For $t, u \in S$ and $X, Y \subseteq S$, we define the *order of T-similitude* $\leq_s \subseteq S \times S$ to be such that

$$t \leq_s u \text{ if and only if } \forall U \subseteq S_0 \mid L_t, L_u \subseteq U.[W_{t,U} \leq^+_{S_0} W_{u,U}]$$

Correspondingly, we define $\leq^\uparrow_s \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ such that

$$X \leq^\uparrow_s Y \text{ if and only if } \forall U \subseteq S_0 \mid L_X, L_Y \subseteq U.[W_{X,U} \leq^+_{S_0} W_{Y,U}]$$

Naturally, if $\mathbf{T}$ is understood, we shall refer to the order of T-similitude as just the order of similitude.

Normally, we would immediately lift $\leq_s$ to a power order in the standard way to derive a relation between subsets of X. For example,

$$X \leq^0_s Y \quad \text{if and only if} \quad \forall t \in X.\exists u \in Y.[t \leq_s u]$$

However, the resulting comparison is incorrect because for $t \in X$ and $u \in Y$, the models of t and u used in the comparison are not necessarily models of X and Y as well. Consequently, in Definition 4.11 we compare X and Y by their models rather than by their members.

Remark 4.12. In the power relation approach to verisimilitude presented in Chapter 4 of [8] (in particular, pp134–138), a propositional language $\mathcal{L}$, freely generated from a countable set of propositional variables, is considered. Each propositional variable is taken to represent an atomic fact. A theory of the language is taken to be a set of propositions, not necessarily deductively closed. A valuation is, as expected, an assignment of true (1) or false (0) to the propositional variables, and valuations are treated as possible worlds. By taking $0 \leq 1$, valuations (and hence possible worlds) can be ordered by taking $u \leq v$ if and only if for any propositional variable p, $u(p) \leq v(p)$ (compare Proposition 2.29). A proposition ϕ is true at a possible world w if the valuation that represents w is a model for ϕ (compare Definition 2.30). The meaning $M(\phi)$ of ϕ is then simply the set of all possible worlds at which ϕ is true, and by extension, the meaning $M(T)$ of a theory T is simply the set of all possible worlds at which every member of T is true. For any two theories T_1 and T_2, the authors define a verisimilar order $\twoheadleftarrow$ by applying the full power order on valuations so that

$$T_1 \twoheadleftarrow T_2 \quad \text{if and only if} \quad \begin{array}{l} \forall w_1 \in M(T_1).\exists w_2 \in M(T_2).[w_1 \leq w_2] \text{ and} \\ \forall w_2 \in M(T_2).\exists w_1 \in M(T_1).[w_1 \leq w_2] \end{array}$$

(see Definition 4.8 in [8]). The approach we have adopted here is not simply a version of verisimilitude that suggests the use of a many-valued logic because in that case also, propositions would be assigned a truth value from a designated set. In our case, the value assigned to a member of X is not restricted to a designated set because the satisfaction reflected by the statement $w \models_v x$ is determined by whether w preserves the value in $\mathfrak{T}_{\mathcal{D}}$ assigned to x by the designated reference assignment v, in the sense of Definition 2.30.

From Definition 4.11, the set $W_{u,U}$ is closer to the idealised state than $W_{t,U}$ if there is a universal increase along $\leq_{S_0}$ in the transition from $W_{t,U}$ to $W_{u,U}$, which we have represented as the full power order $\leq^+_{S_0}$. Consequently, we consider u to represent S_0 more strongly than t, even though possibly $v(u) \leq_T v(t)$. Similarly, for $X, Y \subseteq S$, we understand Y to have a higher degree of similitude or "S_0-likeness" than X if $W_{X,U} \leq^+_{S_0} W_{Y,U}$.

With the order of entrenchment set out in Definition 2.31, sentences in S are effectively compared by the degree to which they are presented by the entity x. The comparisons that we make under $\leq_e$ are more directly relative to x than under $\leq_s$ because assignments now serve as "intermediaries" between the sentences in S and $\leq_e$, whereas with $\leq_s$, assignments determine the order directly. For this reason, under $\leq_e$ we have no ready means of determining the extent to which β is true of x. However, unlike $\leq_s$, $\leq_e$ imposes an order on S that directly reflects how strongly each sentence is presented by x.

Observe that the methods of comparison that we have described here are less direct than the comparisons of Section 4.2.2. There we could state explicitly to what extent β was true of x, whereas here we can do so only by comparing (sets of) sentences. The comparisons made here, however, allow more flexibility because we can account for more complex properties and sets of properties, and hence more complex stereotypes than in Section 4.2.2.

From $\mathbf{T}$, we can once again derive a logic for S as set out in Section 3.2. As described in Section 3.2.3, this logic will reflect $\leq_e$ rather than $\leq_s$ (see Propositions 3.59 and 3.61). Given a family $\{S_{0,i}\}_{i \in I}$ of stereotypes, where each $S_{0,i}$ corresponds to a characteristic β_i, by treating each $\mathbf{S}_i$ as a subject of type $\mathcal{F}$ we can construct a relational B-structure as in Section 4.1.2 to form a unified representation of the stereotypical behaviour of the entity x. In this way, we can reason in a single setting about the different stereotypes presented by x.

Suppose that for each stereotype $S_{0,i}$, x exhibits the properties V_i, of which the stereotypical "consequences" are $C_i(V_i)$, where C_i is the consequence function of the logic derived for $S_{0,i}$. If x begins to present a new property y, any existing contradictory properties should be retired, and in effect a "belief revision" should be conducted on the properties that are presented by x. This revision can be carried out as described in Section 4.1.2.

Finally, given our initial stereotype S_0, suppose we have a stereotype $\overline{S}_0$ that represents the characteristic $\neg\beta$. If we take the reference assignment of x as determining the extent to which β is true of x, then we can form the pair $(v, \overline{v})$, where v and $\overline{v}$ are the reference assignments for S_0 and $\overline{S}_0$ (compare

Section 4.2.2). If we apply the truth order of Section 3.1.2 to these pairs, we then have

$$(v, \overline{v}) \leq (v', \overline{v}') \quad \text{if} \quad v \leq_{S_0} v' \text{ and } \overline{v}' \leq_{S_0} \overline{v}$$

Under this order, the least element $(v_{\mathbf{0}}, \overline{v}_{\mathbf{1}})$ would indicate that $\neg\beta$ was wholly true of x, while the greatest element $(v_{\mathbf{1}}, \overline{v}_{\mathbf{0}})$ would indicate that β was wholly true of x. The remaining values may be taken to represent "in-between" cases where it cannot be said that β or $\neg\beta$ is true of the x.

Summary

In this chapter we set out to develop a belief revision relation in the style of [36] and derived from the dispositions that we modelled in Chapter 3. The belief revision relation that we developed was formulated in terms of belief bases rather than belief sets, because we wished to avoid the computational complexity that accompanies the use of belief sets. Although the relation was set out as a relation between belief sets, it was determined by the properties of the underlying belief bases and in fact also induced a relation between belief bases.

We then constructed a B-structure as a manifold derived from a family of subjects of a specified type. The B-structure that we developed was relational in that each underlying subject was equipped with a relational belief revision operation. To show how belief revision could be conducted with a relational B-structure, we adapted the option B revision scheme of [11] to describe a belief revision operation that was local and relational. In this way, we fulfilled the goal of our dissertation as stated in the abstract and elaborated on in Section 1.3, *viz.* we were able to model belief revision as a local, non-deterministic operation.

Finally, to illustrate the applicability of our work to other research, we recounted the phenomenal, dispositional account of belief presented in [44]. We then provided a worked example in which we showed how the theory we have developed could be applied to this account. To some extent we were able to extend the work of [44] by showing how multiple dispositional stereotypes could be handled and also by showing how stereotypes presented by an entity could be revised if the entity began to exhibit a new stereotypical property.

It is certainly possible to provide a more comprehensive worked example than what we presented, because our purpose was just to illustrate how our work might be applied and to present additional ideas that could be developed further in a future study. Indeed, we have not commented much on how the sheaf theoretic aspects of our model could be applied. In this regard, the papers [49] and [50] provide examples of the benefits of a sheaf-theoretic approach. It is hoped that the short exposition we have provided will serve to generate other ideas, and we shall leave such a development as a possible topic for future research.

Chapter 5

Conclusions

We have now concluded the exposition of our subject matter, a model of belief revision. The central hypothesis of our work has been the idea that

> *Mathematical structures with which to study and translate between global and local behaviour can be applied to model the process of belief revision as a local, non-deterministic operator.*

and through the steps that we set out in Section 1.3 and the framework that we developed over the course of Chapters 2, 3 and 4, we have provided an affirmative response in support of this hypothesis.

In view of the work we have presented, one may immediately ask why we chose to study belief revision in the manner that we did. In fact, we may pose the more general question as to why we should study belief revision at all.

Why should we study belief revision?

Perhaps a simple answer to this question is just "because everything changes". The cornerstones of our world view gradually erode or are reinforced as we adapt and even re-invent ourselves in the process of coping with and managing change. Some of our beliefs are more transitory, such as whether the trains are running on time on a given afternoon, and so are easily relinquished when we find matters contrary to what we thought was the case. Other beliefs are acquired at great cost, and to relinquish them seems somehow to introduce an uncertainty that threatens to unravel the fragile sense of order we have cultivated for ourselves.

Our greater reluctance to relinquish some beliefs than others suggests that our beliefs are not held with equal strength, and in this sense, belief revision can be seen as the guarding of our established beliefs against unnecessary change – the more entrenched a belief, the harder it is to dislodge or replace it. Certainly, however, belief revision is necessary, for we can hardly expect to hold our stock of beliefs constant in the face of all the changes taking place around us – if we were to do so, at some point we would come to hold at least one belief that was contrary to what could be known about our world.

If we agree that belief revision is necessary, perhaps the next question is how the required revision might be effected. To understand this question, we might ask related questions about what sort of changes to our beliefs would constitute a revision, how we would handle information from a possibly unreliable source, or what the role of memory is in the process of revision. Just as the study of logic seeks to make precise the idea of valid inference (compare [12], p1), so the study of belief revision seeks to make precise what constitutes a valid belief revision, and to provide a formal setting in which to explore these kinds of questions.

Arguably, the inevitability of belief revision establishes its significance as well as the importance of its study. Certainly, this importance is attested to by the large body of literature devoted to the subject, some of which we surveyed in Chapter 1. A wide variety of approaches to the problem is represented in this body of literature, which leads us back to the question of why we chose the path that we did.

Our Choice of Approach to Belief Revision.

As described in Section 1.2, the AGM postulates have become the dominating paradigm for reasoning about belief revision. The postulates regulate the action of a belief revision operation, and an operation that satisfies them is then restricted to be a function that maps a deductively closed set to another deductively closed set, where the deductive closure is determined with respect to a given logical language.

Over the course of Section 1.2.1, we described some of the problems induced by the requirement of deductive closure. More specifically, apart from computational complexity, under deductive closure an agent would have to know all the consequences of its beliefs, and all of its beliefs would need to be accounted for during a revision, without regard to the relevance of the beliefs to the new information. Consequently, we chose to work with belief bases rather than belief sets, dispensing with the need for deductive closure.

The choice to work with belief bases introduced the possibility of having a collection of beliefs that was locally consistent but globally inconsistent. To cast a belief revision operation in this setting, we would need to handle this possibility. In this regard, the approaches of [25] and [11] stood apart from other approaches not only in their innovation but also because they addressed the problem of relevance. We chose to adopt the approach of [11] because of its technical simplicity.

The model of [11] is based on an artifice called a B-structure. A B-structure divides a given language into sub-languages, each of which determines a subject area. Each sub-language is then paired with a theory held by an agent about the corresponding subject. The logic is determined by the complete language, rather than being tailored to each subject. Arguably, it may not be appropriate for all the subjects to have the same logic, and in our opinion the innovative model of [11] could be extended by instead adopting a "bottom-up" approach and considering each subject to be autonomous, equipped with its own language, logic and revision operation.

Under certain circumstances, a belief revision operation could have more than one outcome. Consequently, we chose to model belief revision as a relation rather than as a function. To develop a belief revision relation for a subject, we followed the development presented in [36]. We questioned the use of epistemic entrenchments in their approach, however, and chose to extend their work by using the more general notion of a disposition in place of an epistemic entrenchment.

With the belief revision relation in place, we next chose to combine the collection of autonomous subjects by applying the definition, provided in [52], of a manifold in terms of sheaf theory. We could therefore construct a B-structure whose elements were in effect indexed by the given atlas of the underlying manifold, and for which relational belief revision could be conducted both locally within a particular subject and more globally across a combination of subjects. By modelling the process of belief revision as a local, non-deterministic operation, we were then able to fulfil the goal our of dissertation.

Belief revision is a field of study with a vast compass that embraces many complex topics. Consequently, many aspects of the problem lay without the ambit of our work. In particular, one problem that we did not discuss was iterated belief revision, where the beliefs of an agent are revised by some sequence $\varphi_1, \ldots, \varphi_n$ of formulae of a given language. Such revision is straightforward when for any $2 \leq i \leq n$, φ_i is consistent with the beliefs that are held after revising with $\varphi_1, \ldots, \varphi_{i-1}$. Problems arise when instead φ_i is not consistent with these beliefs, for there is at present no clear-cut answer as to how any of the preceding revisions should be "undone" to accommodate φ_i.

We believe that it was to the advantage of our work that we omitted topics such as iterated revision, since their inclusion would have broadened the scope of our exposition to the point where we would have been unable to cover our topic in any adequate depth, rendering our work ineffective. For a detailed critique of approaches to problems in belief revision, the interested reader is respectfully referred to [18] ([19]).

Summary of Contributions and Extensions.

1. *The choice to model belief revision in terms of propositional attitudes.*

This decision marked a departure from the convention of working directly with beliefs, as was done in [1], [11] and [36] for example, and through it we were able to extend the work of [36]. From [45],

> "Contemporary discussions of belief are often embedded in more general discussions of the propositional attitudes; and treatments of the propositional attitudes often take belief as the first and foremost example."

As described in [45], hope, fear, doubt, desire and intention are all examples of propositional attitudes. Using hope as an illustration, an agent m could hope that y was the case more strongly than it did x. Intuitively, a set X of sentences could then be ordered by how strongly m hoped that each $x \in X$ was the case. This similarity to an epistemic entrenchment suggested that hopes could be revised in much the same way as beliefs. Consequently, we chose to formulate our model of belief revision in terms of propositional attitudes in an attempt to generalise and extend the work of [36].

2. *In our development of a B-structure, we chose to use countably rather than finitely many primitives in the construction of each subject.*

If, for a given subject, we had started with a sentence algebra over a non-empty, finite set of primitives, whenever the agent encountered information about the subject that involved new, unknown atoms, it would need to expand this set of primitives and hence also the content of the subject. To stabilise the set of primitives and corresponding subject content, we elected to use countably many primitives instead.

Initially, an agent might know about a finite subset of these primitives only. The remaining primitives would then serve as placeholders for new information that the agent could subsequently learn about. This decision allowed us to extend the work of [11] by including the possibility that an agent might not know about a particular primitive and hence not understand a particular sentence in the content of the subject. We catered for this situation by defaulting any unknown primitives to an indeterminate degree of disposition. In this sense, the agent then implicitly held a degree of disposition towards every primitive and consequently every sentence of the subject content.

The use of countably many primitives exposed us to the risk of having to cater for sentences that contained infinite conjunctions, which the logic of Chapter 3 would not have handled. We excluded these cases by restricting new information to finite sentences in the sense of Definition 4.1. Mechanisms to handle infinite conjunctions lay beyond the scope of our work, and the interested reader is respectfully referred to works such as [5] and [4].

3. *We chose to induce an order on a set of sentences by means of an external component.*

A guiding factor in this choice was the concern that if we imposed the order directly on a given set of sentences, the order might not be invariant under certain transformations. Our work had a predominantly topological setting, and for topological spaces the structure-preserving transformations are homeomorphisms. To illustrate, let $(A, \leq)$ be an ordered set with $A = \{a, b, c, d\}$ and $\leq = \{(a, b), (a, c), (b, d), (c, d)\}$. Given the topological spaces

$$\begin{aligned} A^{\uparrow} &= (A, \{\emptyset, A, \{b, c, d\}, \{b, d\}, \{c, d\}, \{d\}\}) \quad \text{and} \\ A^{\downarrow} &= (A, \{\emptyset, A, \{a, b, c\}, \{b, a\}, \{c, a\}, \{a\}\}) \end{aligned}$$

the map $f : A^{\uparrow} \longrightarrow A^{\downarrow}$ with $f = \{(a, d), (b, c), (c, b), (d, a)\}$ is then an example of a homeomorphism that does not preserve the order on A.

Our primary reason for this choice, however, was the flexibility that we could gain as a result. Our aim was to model the idea that for two sentences x and y, $x \leq y$ if we held a stronger degree of a given disposition towards y than towards x. We developed the idea of a bounded r-algebra of type $\mathcal{F}$ to induce the required order on a given sentence algebra of type $\mathcal{F}$. We could then add as much detail to each bounded r-algebra as was necessary to model a particular order, and because the bounded r-algebras were self-contained, the level of detail in each could vary independently of the other bounded r-algebras.

This independence would also allow us to use different underlying data types for each disposition. For example, we could model one disposition by starting with an enumerated type such as {strongly agree, agree, neutral, disagree, strongly disagree} and deriving surrogate values from that as described in Section 3.1.2, whereas for another disposition we might start with a set of real numbers such as a closed interval and then derive the surrogate values once more. It is not immediately clear that we would have this kind of flexibility at our disposal had we imposed the order on the beliefs directly.

4. *We used a structure called a manifold to combine the subjects into a single, unified structure.*

Two notions at work in a B-structure are unification and autonomy. A sense of autonomy comes about via the decomposition of a theory into subjects that are largely independent, while unification comes about because all of the subjects stem from the same ambient language. Because we began with a family of independently developed subjects, autonomy was present by default. We therefore required a means of accomplishing unification. To this end, we appealed to the definition of a manifold in terms of sheaf theory, as provided in [52], to accomplish the required combination of the autonomous subjects.

5. *The idea, in Definition 2.12, of a local isomorphism of presheaves.*

In this definition, we sought to do for presheaves what local homeomorphisms do for topological spaces. Whereas in Definition 2.37 an equal emphasis is placed on the topology and the sheaf, in Definition 2.12 the emphasis is more on the presheaf than on the topology. Our intention was to allow the exploration of one presheaf in terms of another just as one explores a manifold in terms of other, more well-known spaces.

6. *The idea, in Definition 2.23, of a relational algebraic structure.*

These structures were the primary component of our framework, and combined the idea that a set could carry an algebraic as well as a relational structure. A well-known example is a lattice, where an algebraic structure in the form of meet and join operations has a well-defined interaction with a partial order on a set. For more abstract algebraic types such as those of Chapter 2, the interaction between the algebraic and relational structure on the carrier set would not generally be as well-defined as for lattices. For the work we had to complete, we needed a means of transporting both kinds of structure in a single entity.

7. *The idea, in Definition 2.24, of an r-algebra simulation.*

An r-algebra simulation captures the idea that one r-algebra can simulate the structure of another. The idea is adapted from Definition 4.1.1 in [52], which used the notion of an $\mathbf{R}$-algebra to allow a commutative ring $\mathbf{S}$ with identity to simulate another such ring $\mathbf{R}$. A sheaf of r-algebra simulations gives us a simple means of equipping a topological space with the structure of a given r-algebra.

In our case, the simulations were accomplished with homomorphisms. However, other forms of simulation are available as well. For example, from [2] we have the idea of a bounded morphism. For any two r-algebras $\mathbf{A}$ and $\mathbf{B}$ of type $\mathcal{F}$, we might then define $f : A \longrightarrow B$ to be a bounded morphism from $\mathbf{A}$ to $\mathbf{B}$ if f is a homomorphism from (A, F) to (B, F), and if $f(a)R_B y$ then there is $b \in A$ with $aR_A b$ and $f(b) = y$.

8. *The order of entrenchment (Definition 2.31) and the order of similitude (Definition 4.11).*

The order of similitude is an adaptation of the work on pp134–138 in Chapter 4 of [8]. We applied the idea of verisimilitude via power relations to the idea of dispositions, so that the reference assignment of an agent might approach the idealised state of holding the maximum degree of disposition towards all the primitives of a given language.

For both orders, we used the idea from [8] that the meaning of a particular sentence can be determined from its models. In [8], these models were simply those valuations that assigned the sentence the value true. By treating valuations as states of affairs, the meaning of a sentence was then simply those states of affairs in which the sentence was true. In contrast, we described in Definition 2.30 a notion of satisfaction based on the idea of preservation of degrees of truth as set out in [16]. With this idea, a reference valuation v gave rise to a set T_v of designated degrees of disposition towards a sentence x, and a valuation w was a model of x if $w(x) \in T_v$.

For the order of similitude, we ordered sentences by considering their meanings directly, using valuations as indicators of the strength of a particular sentence with respect to the given disposition. In contrast, for the order of entrenchment the valuations served as intermediaries between sentences and the order, and we ordered sentences according to the values assigned to them by a common set of states of affairs. This gave rise to an order that was similar in character to that of [36].

The (surrogate) degree of disposition assigned to a sentence was computed by means of the reference valuation. For any sentence in the given sentence algebra, the degree of disposition could then be computed based on the values assigned to the primitives of the language. We consider this an advantage over the use of epistemic entrenchments, where the assignments are assumed and "pre-existing", rather than computed. Furthermore, for any new sentence we could determine the location of the sentence in the order, whereas this was not generally possible for epistemic entrenchments until such a value had been decided on.

9. *The ordered manifolds of type* M *over an r-algebra (Definition 2.39) and the construction procedure of Section 2.3.2.*

These manifolds were an adaptation to the case of r-algebra simulations of the manifolds described by Definition 4.3.6 in [52]. They provided a means whereby an ordered space over a given r-algebra could be constructed from a family of ordered spaces over the given r-algebra.

To construct such a manifold, we adapted the approach set out in Lemma 1.23 on pp21-22 of [33]. The approach of [33] was simpler than in our case because it was based on a topological n-manifold for which the (optional) structure was conferred without recourse to sheaf theory. We extended this approach to allow a particular structure to be conferred on the manifold by a structure sheaf, and in the process, we could construct a manifold of the kind described by Definition 4.3.6 in [52]. The properties and local structure of the manifold could then be explored and described in terms of the family of constituent spaces by using an atlas of charts, similarly to [33].

10. *The model of disposition of Chapter 3.*

This model is one of the main contributions of our work. To our knowledge, a model such as ours has not been developed, at least not with the same goal in mind. Our model of a disposition as a partially ordered set of surrogate values was derived purely from a family of observations about a set of values whose exact nature it was not necessary to know. Our hope is that if not the structure, then at least the technique of derivation will be useful in future research in the field of belief revision or other fields of study such as program semantics.

In developing the model, we were able to extend the work of [22] and [23] in that we first used the given observations to (implicitly) order the points in our set of degrees of disposition, then lifted the order to a power order so we could order the observations.

The relation $\preceq$ of Definition 3.6 may equivalently be written as

$$c \preceq d \quad \text{if and only if} \quad \forall U \in \mathcal{D}.[c \in U \Rightarrow d \in U]$$

This form of the definition is based on that presented in [22] and [23]. In both of those works however, the antecedent and consequent on the right are interchanged so that

$$c \preceq d \quad \text{if and only if} \quad \forall U \in \mathcal{D}.[d \in U \Rightarrow c \in U]$$

The interchange is motivated by the notion of topological convergence, and c may be understood as providing approximate information about d by virtue of the observations that c and d both satisfy. The point d then serves as an idealised limit to the approximate information provided by c. In contrast, it was the notion of refinement that interested us, and in Definition 3.6 the point d was a refinement of c because d satisfied more observations than did c and so represented a stronger property than c, as shown by Proposition 3.9.

The idea of associating a point with a set of properties is not new. For example, in [41] this kind of association is represented as a relation of type $X \times \mathcal{P}(X)$ on a set X. These relations are called binary multi-relations and have found application particularly in the field of program semantics.

The idea is that a given input state is mapped by a binary multi-relation to a set of properties that the output from a particular program can be expected to satisfy. The binary multi-relation is then used to represent the program and captures in an accessible form angelic or demonic non-deterministic terminating behaviour of the program.

In our case, we have turned this picture around by instead using properties to determine points, thereby focussing on properties of points rather than points themselves. The idea of using observations to structure the points of a set is reminiscent of the work presented in [54], where the "opens" of a frame are used for this purpose. Scope did not permit us to explore the development of relational belief revision within settings such as those of [54]. It is our belief, however, that such exploration would be valuable research because techniques from locale theory and formal (point-free) topology could then be applied to the problem of belief revision.

In deriving our set of surrogates, we based our work on [15]. There were, however, some important differences. In [15], a topology was assumed, and truth values were composed from the open and closed sets of the topology. A truth value was thus constructed from two properties U and V, which respectively were thought of as sets of worlds at which a formula of interest held and did not hold. We extended this idea by composing a surrogate from two sets of properties. We could therefore uncouple surrogate values from direct association with the underlying degrees of disposition.

To form a surrogate, we chose to use the closed sets of a closure operator (compare the discussion following Proposition 3.11). As with [15], the resulting lattice of surrogates was finitely distributive. In our case, however, we could define negation as an involution (Definition 3.18 and Proposition 3.19), whereas in [15], negation corresponded to the weaker notion of proto-complementation (see Definition 5.14 in [12]).

From our development of a model of disposition, there arise several avenues for such future research. It would be of interest to explore how r-algebra simulations might be applied to other fields of study. For example, one very simple application to the study of program semantics might be as follows. Let $\mathbf{A}$ be a bounded r-algebra of type $\mathcal{F}$, where $(A, \leq_A)$ represents a partially ordered family of properties, and let $(S, \Omega S)$ be the state space for a given program. For each $U \in \Omega S$, we let $[U \longrightarrow A]$ be the family of functions such that if $f(u) = \theta$ for $f \in [U \longrightarrow A]$, then when the program is started in state u, the outcome satisfies the property θ.

From Proposition 2.29, the sheaf H on S is such that for each $U \in \Omega S$, $H(U) = [U \longrightarrow A]$, is a sheaf of $\mathbf{A}$-simulations. We can then represent the actual behaviour of the program by a family $\{s_i\}_{i\in I} \subseteq H(S)$ of global sections. With restriction maps, the program behaviour can be analysed by considering more local behaviour over an open set of interest. We may also think of $\{s_i\}_{i\in I}$ as representing a relation $R \subseteq S \times A$ where for $x \in S$, $R(s) = \{s_i(x)\}_{i\in I}$. We are then able to represent a form of non-determinism, from which we can proceed to explore other aspects such as the non-terminating behaviour of a program.

In turn, a manifold would provide a means of combining the state spaces and behaviour of several programs into a unified structure. The behaviour of each program can be examined by means of the accompanying atlas, and where two programs share input states, the transition maps enforce a coherency on the behaviour of the programs on these shared inputs. If each program is treated as a method of a class in an object-orientated setting, we can then reason about the behaviour of the class. Similarly, if each program is treated as an independent agent, we may reason about the behaviour of a family of co-working agents.

We could use the P- and Q-specifications of a surrogate to specify, respectively, "liveness" (something good always happens) and "safety" (something bad does not happen) properties (see for example [47] and also [3]). An overdefined truth value assigned to an input state is then more desirable, since for output states that satisfy P and Q something good always happens and something bad does not happen. Consistent and exact truth values are less desirable, because for output states that satisfy P, we cannot guarantee that something bad does not happen. It may also be of interest to explore how these surrogates could be used in conjunction with other analysis tools such as predicate transformers.

In closing our work, we recall again the quote from [29] (page XX):

> "Abstract algebra cannot develop to its fullest extent without the infusion of topological ideas, and conversely, if we do not recognise the algebraic aspects of the fundamental structures of analysis our view of them will be one-sided."

Although we did not need to explore the kind of duality suggested by this quotation, we have nonetheless been fortunate to have explored a problem in belief revision from logical, algebraic and topological vantage points. The logical view would always be a given, because belief revision is so intrinsically linked to the idea of logical consistency. Ideas from universal algebra were applied most notably in the exposition of the Chapter 2 and to a lesser extent in the development of the dispositions of Chapter 3. It was ultimately the application of topological ideas, however, that allowed us to develop our proposed framework.

It is our hope that the framework we have constructed will be of benefit to further study of belief revision, and that it may find application in other areas of study as well.

Bibliography

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, Cambridge, UK, 1st edition, 2001.

[3] Marcello Bonsangue. *Topological Dualities in Semantics*. PhD thesis, Institute for Programming Research and Algorithmics, Vrije Universiteit Amsterdam, 1996.

[4] Marcello M. Bonsangue and Joost N. Kok. Towards an Infinitary Logic of Domains: Abramsky Logic for Transition Systems. *Information and Computation*, 155:170–201, 1999.

[5] Marcello M. Bonsangue and Joost N. Kok. Infinite Intersection Types. *Information and Computation*, 186(2):285–318, 2003.

[6] Craig Boutilier. A Unified Model of Qualitative Belief Change: A Dynamical Systems Perspective. *Artificial Intelligence*, 98(1–2):281–316, 1998.

[7] Craig Boutilier, Ronen I. Brafman, Carmel Domschlak, Holger H. Hoos, and David Poole. CP-nets: A Tool for Representing and Reasoning with Conditional Ceteris Paribus Preference Statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2003.

[8] Chris Brink and Ingrid Rewitzky. *A Paradigm for Program Semantics: Power Structures and Duality*. Number 17 in Studies in Logic, Language and Information. CSLI Publications, Stanford University, USA, 1st edition, 2001.

[9] Lawrence M. Brown. Quotients of Textures and of Ditopological Texture Spaces. In J. Dydak, G. Gruenhage, J. Heath, et al., editors, *Topology Proceedings*, volume 29(2), pages 337–368. Nipissing University, Ontario, Canada, 2005.

[10] Stanley Burris and H. P. Sankappanavar. *A Course in Universal Algebra: The Millenium Edition*. Springer Verlag, 2000.

[11] Samir Chopra and Rohit Parikh. Relevance Sensitive Belief Structures. *Annals of Mathematics and Artificial Intelligence*, 28(1–4):259–285, 2000.

[12] John P. Cleave. *A Study of Logics*. Oxford University Press, New York, 1st edition, 1991.

[13] B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, Cambridge, UK, 2nd edition, 2002.

[14] Renée Elio. What to Believe when Inferences are Contradicted: The Impact of Knowledge Type and Inference Rule. In Michael G. Shafto and Pat Langley, editors, *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society, Stanford University Aug 7-10 1997*, pages 211–216. Lawrence Erlbaum Associates, New Jersey, USA, 1997.

[15] Melvin Fitting. Logic Programming on a Topological Bilattice. *Fundamenta Informaticae*, XI:209–218, 1988.

[16] Josep Maria Font. An Abstract Algebraic Logic View of Some Multiple-Valued Logics. In Melvin Fitting and Ewa Orlowska, editors, *Beyond Two: Theory and Applications of Multiple-Valued Logic*, Studies in Fuzziness and Soft Computing, pages 25–57. Physica-Verlag, Heidelberg, Germany, 2003.

[17] Josep Maria Font, Ramon Jansana, and Don Pigozzi. A Survey of Abstract Algebraic Logic. *Studia Logica, Special Issue on Algebraic Logic, Part II*, 74:13–97, 2003.

[18] Nir Friedman and Joseph Y. Halpern. Belief Revision: A Critique. In Luigia Carlucci Aiello, Jon Doyle, and Stuart Shapiro, editors, *KR'96: Principles of Knowledge Representation and Reasoning*, pages 421–431. Morgan Kaufmann, San Francisco, California, 1996.

[19] Nir Friedman and Joseph Y. Halpern. Belief Revision: A Critique. *Journal of Logic, Language and Information*, 8(4):401–420, 1999.

[20] Peter Gärdenfors and David Makinson. Revisions of Knowledge Systems using Epistemic Entrenchment. In Moshe Y. Vardi, editor, *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, California, USA, March 1988. Morgan-Kaufmann.

[21] Edmund L. Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.

[22] Jonathan Gratus and Timothy Porter. A Geometry of Information, I: Nerves, Posets and Differential Forms. In Ralph Kopperman, Michael B. Smyth, Dieter Spreen, and Julian Webster, editors, *Spatial Representation: Discrete vs. Continuous Computational Models*, number 04351 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2005. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany.

[23] Jonathan Gratus and Timothy Porter. A Geometry of Information, II: Sorkin Models, and Biextensional Collapses. In Ralph Kopperman, Michael B. Smyth, Dieter Spreen, and Julian Webster, editors, *Spatial Representation: Discrete vs. Continuous Computational Models*, number 04351 in Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Dagstuhl, Germany, 2005.

[24] Thomas D. Griffin and Stellan Ohlsson. Beliefs versus Knowledge: A Necessary Distinction for explaining, predicting, and assessing Conceptual Change. In Johanna D. Moore and Keith Stenning, editors, *Proceedings of the 23rd Annual Conference of the Cognitive Science Society 1-4 August 2001, University of Edinburgh, Scotland*, pages 364–369, New Jersey, USA, 2001. Lawrence Erlbaum Associates, Inc.

[25] Sven Ove Hansson and Renata Wassermann. Local change. *Studia Logica*, 70(1):49–76, 2002.

[26] Andreas Herzig and Omar Rifi. Propositional Belief Base Update and Minimal Change. *Artificial Intelligence*, 115(1):107–138, 1999.

[27] John H. Holland. *Emergence: From Chaos to Order*. Oxford University Press, Oxford, UK, 1st edition, 1998 (First paperback edition 2000).

[28] Kiyosi Itô, editor. *Encyclopedic Dictionary of Mathematics*. MIT Press, Cambridge, MA, USA, 2000 (5th Printing).

[29] Peter T. Johnstone. *Stone Spaces*, volume 3 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, UK, 1982 (reprinted 1992).

[30] Hirofumi Katsuno and Alberto Mendelzon. On the Difference Between Updating a Knowledge Base and Revising It. In James F. Allen, Richard Fikes, and Erik Sandewall, editors, *KR'91: Principles of Knowledge Representation and Reasoning*, pages 387–394. Morgan Kaufmann, San Mateo, California, 1991.

[31] Hirofumi Katsuno and Alberto Mendelzon. On the Difference Between Updating a Knowledge Base and Revising It. In Peter Gärdenfors, editor, *Belief Revision*, number 29 in Cambridge Tracts in Theoretical Computer Science, pages 183–203. Cambridge University Press, Cambridge, UK, 1992.

[32] Jérôme Lang. Belief Update Revisited. In Manuela M. Veloso, editor, *IJCAI'07: Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6–12, 2007*, pages 2517–2522. IJCAI, and AAAI Press, Menlo Park, California, 2007.

[33] John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer Science+Business Media, Inc., New York, USA, 1st edition, 2003.

[34] John M. Lee. *Introduction to Topological Manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer Science+Business Media, LLC., New York, USA, 1st edition, 2004.

[35] Isaac Levi. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press, USA, 1980 (First paperback edition 1983).

[36] Sten Lindström and Wlodzimierz Rabinowicz. Epistemic Entrenchment with Incomparabilities and Relational Belief Revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change (Workshop, Konstanz, FRG, October 13–15, 1989 Proceedings)*, volume 465 of *Lecture Notes in Artificial Intelligence (Subseries of Lecture Notes in Computer Science)*, pages 93–126. Springer Berlin/Heidelberg, Germany, 1991.

[37] Saunders MacLane. *Categories for the Working Mathematician*, volume 5 of *Graduate Texts in Mathematics*. Springer-Verlag New York, Inc., USA, 2nd edition, 1978.

[38] Saunders MacLane and Ieke Moerdijk. *Sheaves in Geometry and Logic - A First Introduction to Topos Theory*. Springer-Verlag New York, Inc., USA, 1992.

[39] Rohit Parikh. Beliefs, Belief Revision, and Splitting Languages. In Lawrence Moss, Jonathan Ginzburg, and Maarten de Rijke, editors, *Logic, Language and Computation*, volume 2(96) of *CSLI Lecture Notes*, pages 266–268. CSLI Publications, 1999.

[40] Timothy Porter. Can Categorical Shape Theory Handle Grey-level Images? In Ying-Lie O, Alexander Toet, David Foster, Henk J.A.M. Heijmans, and Peter Meer, editors, *Shape in Picture: Mathematical Description of Shape in Grey-level Images (Proceedings of the NATO Advanced Research Workshop "Shape in Picture", Driebergen, The Netherlands, 7–11 September 1992)*, volume 126 of *NATO ASI Series F:Computer and Systems Sciences*, pages 127–146, Heidelberg, Germany, 1994. Springer-Verlag.

[41] Ingrid Rewitzky. Binary Multirelations. In Harrie C.M. de Swardt, Ewa Orlowska, Gunther Schmidt, and Marc Roubens, editors, *Theory and Applications of Relational Structures as Knowledge Instruments*, volume 2929 of *Lecture Notes in Computer Science*, pages 256–271. Springer, 2003.

[42] Hans Rott. Two Dogmas of Belief Revision. *The Journal of Philosophy*, 97(9):503–522, 2000.

[43] Kenneth M. Sayre. *Belief and Knowledge: Mapping the Cognitive Landscape*. Rowman and Littlefield Publishers, Inc., Maryland, USA, 1997.

[44] Eric Schwitzgebel. A Phenomenal, Dispositional Account of Belief. *Noûs*, 36(2):249–275, 2002.

[45] Eric Schwitzgebel. Belief. In Edward N. Zalta, editor, *The Stanford Encyclopaedia of Philosophy*. Fall 2008 edition, 2008.

[46] Philippe Smets and Robert Kennes. The Transferable Belief Model. *Artificial Intelligence*, 66(2):191–234, 1994.

[47] Michael B. Smyth. Topology. In S. Abramsky, D. Gabbay, and T.S.E. Maibaum, editors, *Mathematical Structures*, volume 1 of *Handbook of Logic in Computer Science*, New York, USA, 1992. Oxford University Press.

[48] Viorica Sofronie-Stokkermans. *Fibred Structures and Applications to Automated Theorem Proving in Certain Classes of Finitely-Valued Logics and to Modeling Interacting Systems*. PhD thesis, Forschungsinstitut für Symbolisches Rechnen: Technisch-Naturwissenschaftliche Fakultät, Johannes Kepler University, Linz, 1997.

[49] Viorica Sofronie-Stokkermans. Sheaves and Geometric Logic and Applications to Modular Verification of Complex Systems. volume 230 of *Electronic Notes in Theoretical Computer Science*, pages 161–187. 24 March 2009.

[50] Viorica Sofronie-Stokkermans and Karel Stokkermans. Modeling Interaction by Sheaves and Geometric Logic. In G. Ciobanu and Gh. Paun, editors, *Proceedings of FCT'99*, volume 1684 of *Iasi, LNCS*, pages 512–523. Springer, 31 August – 3 September 1999.

[51] Matthias Steup. The Analysis of Knowledge. In Edward N. Zalta, editor, *The Stanford Encyclopaedia of Philosophy*. Fall 2008 edition, 2008.

[52] B.R. Tennison. *Sheaf Theory*, volume 20 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, UK, 1975 (Re-issued 2007).

[53] David Truncellito. Epistemology. In James Fieser and Bradley Dowden, editors, *Internet Encyclopedia of Philosophy*. 2007 (Downloaded 27 June 2010). http://www.iep.utm.edu/epistemo.

[54] Steven Vickers. *Topology via Logic*, volume 5 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, UK, 1st edition, 1989.

[55] Luminiţa Vîţă and Douglas S. Bridges. A Constructive Theory of Point-Set Nearness. *Theoretical Computer Science*, 305(1–3):473–489, 2003.

[56] Stephen Willard. *General Topology*. First published by Addison Wesley (1970), republished by Dover Publications (2004), New York, USA, 2004.