Reasoning with Inconsistent Information

Paul Wong

A thesis submitted for the degree of
Doctor of Philosophy at
The Australian National University

December 2004

Typeset in Palatino and Euler by $\TeX$ and $\LaTeX\,2_\varepsilon$.

Except where otherwise indicated, this thesis is my own original work.

Paul Wong
1 December 2004

To Fiona, Little Woks, Smallness, Tickles, Tigger, Gollie and Mother Superior.

Acknowledgements

My undying gratitude goes to all those who helped me along the way – members of my committee, Dr. John Slaney, Dr. Rajeev Gore, Professor John Lloyd, and Professor Robert Meyer for their wise advice and for providing me with the intellectual freedom to explore my own ideas, Dr. Philippe Besnard for his patience and willingness to listen to my absurd ideas and co-authorship of papers (and thereby assisting me in acquiring my first Erdös number). I am also indebted to Professor Ray Jennings and Professor Peter Schotch for their inspiration, and to Professor Peter Appostoli and Professor Bryson Brown for their elegant proof of the completeness of K_n. I have also benefited from conversations with Professor Graham Priest, Professor Chris Mortensen, Professor Dov Gabbay and Professor John Wood about all things inconsistent. My examiners, Dr. Anthony Hunter, Professor Greg Restall and Dr. Edwin Mares also deserve special thanks for their careful and generous comments.

My thanks also go to Dr. Timothy Surendonk and Dr. Glenn Moy for their friendship and assistance in settling into Canberra initially. I am also thankful to my logic siblings David Low and Dr. Ng Ping Wong from 'up above' and Nicolette Bonnette, Kalhil Hodgson, and Dr. Andrew Slater from 'down under'. Nicholas and Chwee Chwee von Sanden offered their generous friendship while I was an apprentice lecturer at the University of Wollongong during the academic year of 2002-2003. Thanks also go to other CSL suspects especially Dr. Jen Davoren for her encouragement and cool hair parties, Dr. Matthias Fuchs for being the impeccable lunch police and for his great passes during soccer, Dr. Doug Aberdeen for starting the construction of the Coke Castle (and thereby provided the motivation for the exponential consumption of cokes in the lab), Dr. Arthur Gretton for starting the Star Wars collectors' card competition. I thank all the CSL students who provided endless hours of amusement, laughter, and excuses for procrastination – Evan, Edward, Kee Siong, Cheng, Agnes, Phil, Dave, Kerry, Vaughan, Pietro, Charles, Greg, David, Jemma and Tatiana.

I would like to thank the Commonwealth of Australia for providing an IPRS Scholarship and the Australian National University for providing an ANU scholarship.

. . . [M]ysticism might be characterized as the study of those propositions which are equivalent to their own negations. The Western point of view is that the class of all such propositions is empty. The Eastern point of view is that this class is empty if and only if it isn’t.

Raymond Smullyan [174]

Abstract

In this thesis we are concerned with developing formal and representational mechanisms for reasoning with inconsistent information. Strictly speaking there are two conceptually distinct senses in which we are interested in reasoning with inconsistent information. In one sense, we are interested in using logical deduction to draw inferences in a symbolic system. More specifically, we are interested in mechanisms that can continue to perform deduction in a *reasonable* manner despite the threat of inconsistencies as a direct result of errors or misrepresentations. So in this sense we are interested in inconsistency-tolerant or paraconsistent deduction.

However, not every case of inconsistent description is a case of misrepresentation. In many practical situations, logically inconsistent descriptions may be deployed as representations for problems that are inherently conflicting. The issue of error or misrepresentation is irrelevant in these cases. Rather the main concern in these cases is to provide meaningful analyses of the underlying structure and properties of our logical representation which in turn informs us about the salient features of the problem under consideration. So in this second sense, we are interested in deploying logic as a *representation* to model situations involving conflict.

In this thesis we adopt a novel framework to unify both logic-as-deduction and logic-as-representation approaches to reasoning with inconsistent information. From a *preservational* view point, we take deduction as a process by which metalogical properties are preserved from premises to conclusions. Thus methodologically we may begin by identifying inconsistency-tolerant deduction mechanisms and then investigate what additional properties of inconsistent premises are preserved by these mechanisms; or alternatively we may begin by identifying properties of inconsistent logical descriptions and investigate which deductive mechanisms can preserve these properties. We view these as two aspects of the same investigation. A key assumption in this work is that adequate analyses of inconsistencies require provisions to quantitatively measure and compare inconsistent logical representations. While paraconsistent logics have enjoyed considerable success in recent years, proper quantitative analysis of inconsistencies seems to have lapsed behind to some extent. In this thesis we'll explore different ways in which we can compare and measure inconsistencies. We hope

to show that both inference and analysis can fruitfully be brought to bear on the issue of inconsistency handling under the same methodological scheme.

Contents

Chapter 1

Introduction

1.1 Motivation

It is customary nowadays to begin a thesis with some remarks about the motivation behind the work. This is typically done with the aid of an example. This thesis is no exception. We shall begin with the following imaginary scenario which we shall call the information fusion problem. Consider a situation in which an object O may be located in one out of nine distinct possible locations represented by a 3×3 grid. Information about the location of O is encoded in a simple propositional language with p's representing the rows and q's representing the columns (see figure (1.1)). Furthermore, complex expressions are generated using the usual Boolean connectives $\{\neg, \wedge, \vee\}$ with their usual truth conditions. We are interested in locating O, and information is gathered from various sensors or sources about the location of O.

	q_1	q_2	q_3
p_1	$\times$		$\times$
p_2			
p_3			

Figure 1.1: A simple logical representation of an object's location.

Suppose we receive two messages:

$$A : p_1 \qquad B : \neg q_2$$

From the received messages we conclude that the possible location of O is:

$$C : (p_1 \wedge q_1) \vee (p_1 \wedge q_3)$$

Our example highlights several important methodological points. The first is the obvious point that information about the state of the world can be encapsulated in a formal language. The practical corollary of this is that more expressive formal languages are required for more demanding representational tasks. But more importantly, since a more expressive language may involve a greater computational cost, the choice of language should be gauged in terms of the representational task at issue. In our example it is clear that a simple propositional language suffices for the representational task.

The second point is that contextual information is often crucial to a reasoning task. In our example, the background information is that the object O is located in exactly one and no more then one location, and that there are exactly nine possible locations of O. It is only in the context of this background information that we can deduce C from A and B. More importantly, background information is not always explicitly stated in a given situation.

Thirdly, our example illustrates how the process of reasoning can be viewed as exploration in the space of possibilities – eliminating some and further exploring others. Our symbolic representations A and B impose certain restrictions on the space of possibilities. These expressions have *truth conditions* which inform us that the world is one way but not another. Furthermore, this information is compositional in the sense that the aggregate of A and B is simply the aggregate of their restrictions on the space of possibilities. The conclusion C is simply what is possible relative to the restrictions imposed by A and B together with the background information.

Finally, our example also illustrates the role and importance of *uncertainties* in reasoning. If we want to know whether O is located at $p_1 \wedge q_1$, the information given is *insufficient* to answer our question. In this sense, the given information is incomplete with respect to our query. Now just as it is possible that we may have incomplete information, it is equally possible that we may have *too much* information – we may receive a third message:

$$D : p_3$$

D is not consistent with A since our background assumption is that the nine locations are distinct and no physical object can be at different places at the same time. In short, there is no guarantee that the information we gather from different sources is either complete or consistent. The possibility of misrepresentation or error is a genuine threat in the process of information fusion.

1.2 Two Approaches to Inconsistencies

As the title suggests, the purpose of this thesis is to examine and develop formal and representational mechanisms for reasoning with inconsistent information. Strictly speaking there are at least two conceptually distinct senses in which we are interested in reasoning with inconsistent information. In one sense, we are interested in using logical deduction to draw inferences in a symbolic system. This is the traditional approach to AI where a logic is used to perform deduction over a logical description or a *knowledge base* representing various states of the environment external to the system. But we are interested in more then that. As we have illustrated with the problem of information fusion, we are interested in mechanisms that can continue to perform deduction in a *reasonable* manner despite the threat of inconsistencies. In these cases, our logical descriptions are inconsistent essentially because they have misrepresented the external environment. So in this sense our deduction must be fault tolerant – it must operate under the explicit assumption that the input data may be erroneous or unreliable. However, not every case of inconsistent description is a case of misrepresentation. Consider for instance,

- negotiation amongst agents with conflicting goals, e.g. selling at the highest price vs buying at the lowest price
- constitutions or legal documents in which incompatible rules apply to the same situation, e.g. you must obey the speed limit vs you must maintain a speed that is consistent with the other traffic
- software requirements engineering process in which different stake-holders have different and incompatible requirements, e.g. ease of use vs advanced features
- faulty artifacts and systems in which expected behaviours diverge from their observed behaviours, e.g. brake lights should be on when braking occurs vs no brake lights when braking
- constraint satisfaction problems that are over-determined, e.g. no one should work more then 8 hours in any given work day vs there is a shortage of staff to cover all work days

In these cases, our logical descriptions may be inconsistent because they correctly represent situations or problems involving conflict in one form or another. So there need not be any misrepresentation involved. The main issue in these cases is not fault tolerant deduction per se but an analysis of the structure and the underlying properties of our logical representation which in turn informs us about the nature of the situation or problem under consideration. Of course to provide such an analysis, our logical description must capture the salient features of the problem *at some appropriate level of abstraction*. But this is very much a question of the representational efficacy of the formal language and not so much a question about deduction. So in another sense, we are interested in deploying logic as a *representation* to model situations and problems involving conflict.

In terms of using a formal language to model real world problems, there is a subtle question as to whether inconsistency in the strictly logical (proof theoretic or model theoretic) sense is the right formalism for modelling conflicts. Conceivably, we can deploy a very different formalism to represent these problems so that the resulting representation is no longer inconsistent in the strictly logical sense. But to do so is to miss an important point. What makes an over-constrained scheduling problem interesting and difficult is precisely that the real world cannot meet its demands. A change in formalism may allow us to find hidden structures of the problem more easily or to perform computation over the representation more efficiently, but this by itself would not resolve the underlying conflict. There is a genuine sense in which the salient feature of conflict is captured in terms of logical inconsistency.

So a map of our conceptual space should include at least two distinct senses of 'reasoning with inconsistent information':

	Inconsistency as error	Inconsistency as conflict
Logical deduction as inference	$\surd$	
Logical description as representation		$\surd$

Figure 1.2: Reasoning with inconsistent information: a conceptual map

Underlying these different senses of 'reasoning with inconsistent information' is of course the traditional distinction between the proof theory and the model theory of a logic. From a purely theoretical standpoint, these are of course distinct and independent approaches to studying logic. But from a computational and system design standpoint, it makes good sense to consider a single symbolic system that can incorporate both the functionality of a prover and the functionality of an analyst for handling inconsistent information. Viewed as a fault tolerant reasoner our prover should provide support for drawing inferences from inconsistent and erroneous data. But viewed as a modeller of problems our analyst should provide support for extracting useful information and patterns from data that represent situations involving conflicts. The overall architecture of such a hybrid system is depicted in figure (1.3).

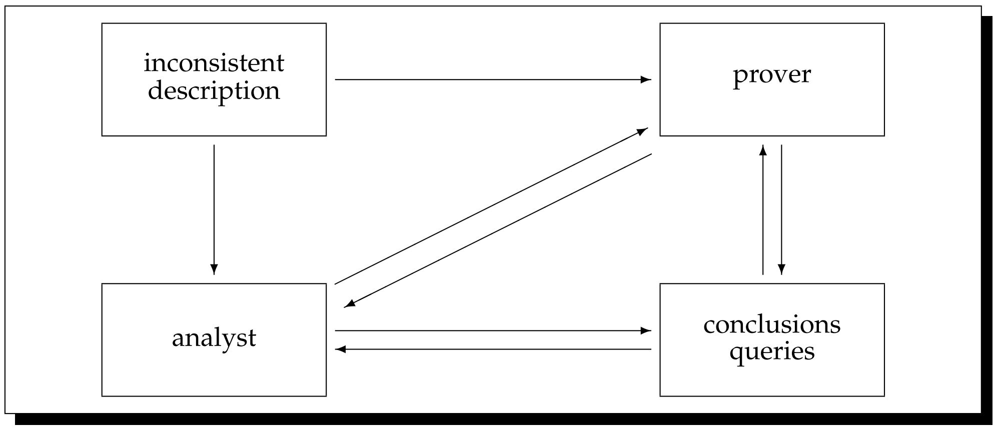

Figure 1.3: A hybrid symbolic system for handling inconsistent information

Indeed, the idea of such a hybrid system is not new. In [172; 173], Slaney and co-workers have proposed and implemented the system **SCOTT** (Semantically Constrained **OTTER**) and more recently **MSCOTT** (Multi-**SCOTT**) which combines the first-order resolution prover **OTTER** with the finite model generator **FINDER** (Finite Domain Enumerator). For our purposes here, we need not be concerned about the detailed workings of **SCOTT**. Figure (1.3) doesn't in fact capture the structure of **SCOTT** – for instance, **FINDER** and **OTTER** do not communicate directly with each other. The main point is that **SCOTT** is a system that combines both reasoning and modelling to accomplish its task. The overall philosophy behind **SCOTT** is to inject some intelligence into a prover by providing semantic information to assist in proof search.

Stated as such, the aim of **SCOTT** is clearly directed towards theorem proving. So in this respect, the role of **FINDER** is mainly to provide assistance to **OTTER**. But we need not think of our hybrid system merely as a theorem proving system. For our purposes, it would be more advantageous to view such a system as a *practical reasoning system*. Firstly, it is practical in the sense that, unlike **SCOTT**, its target domain of application need not be limited to proving mathematical theorems. Like many traditional knowledge base systems such as relational or deductive databases, we can view our hybrid system as a symbolic system for representing and reasoning with information about the external environment. So in this sense, the knowledge base of the system may include *empirical* information of various sorts. Secondly, our hybrid system is also practical in the sense that it is goal directed. The goal of the system may be specified by its immediate user, or alternatively the system may operate as a component of a larger complex system. In either case, what our hybrid system does with its knowledge base depends on what the user or the rest of the system wants. This view of our hybrid system is of course more open-ended. But it does capture both senses of reasoning with inconsistent information within the framework of a single system. As a fault tolerant reasoning system for handling information fusion, the emphasis is perhaps on deduction where the role of the analyst is to assist the prover. But as a modelling system for representing problems or situations, the emphasis is perhaps on analysis where the role of the prover is to assist the analyst. In this thesis, our aim is to explore the theoretical foundation for such a system. In particular we would like to offer a novel theoretical underpinning of the interaction between the analyst and the prover. We would like to consider how inference and analysis can fruitfully be brought to bear on the issue of inconsistency handling under one and the same conceptual scheme.

1.3 Symbolic and Numeric Approaches to Uncertainty

In a broader context, the problem of reasoning with inconsistent information is of course a special case of the more general problem of reasoning with uncertainties. Like reasoning with inconsistent information, reasoning with uncertain information can be viewed as an inference problem as well as a representation problem. In recent years, AI researchers have focused on two general approaches to uncertainty which closely parallel our interests in the use of logical deduction to draw inferences and the use of logical description as a representation. In the traditional *symbolic* approach the emphasis is on developing logical mechanisms for handling uncertainty. This includes developing new deduction mechanisms together with corresponding semantics for reasoning with uncertain information. The symbolic approach has a long history dating back to the works of Newell and Simon on General Problem Solver [133], McCarthy's works on circumscription [128], Reiter's works on Default logics [151], Doyle's works on Truth Maintenance Systems [63] and de Kleer's extension to Assumption-based Truth Maintenance Systems [54; 55]. In contrast, the emphasis of the more recent *numeric* approach is on the representation of uncertainty. Here, the key concern is to develop quantitative methods for measuring uncertainties. This includes approaches that are based on complete ordering, e.g. fuzzy set theory and its offspring possibility theory, as well as approaches that are based on counting, e.g. statistical or probabilistic methods (for a review see chapter 1 and 2 of [95]). In this thesis we would like to consider the symbolic and numeric approaches to inconsistency within the general framework of our hybrid symbolic system. Indeed, a main assumption behind this thesis is that these two approaches should not be viewed as competing strategies for managing inconsistency. On the one hand, there is no doubt that we need to develop inference mechanisms that can perform deduction in a principled way in an environment in which inconsistencies may arise. The underlying assumption is that on occasions it may be desirable to *tolerate* the presence of inconsistencies rather than *revising* one's data. This is especially important in situations in which the turn-over rate of information is much higher than the rate at which consistency checks can be made. On the other hand, there is also a need to develop theoretical and conceptual tools to analyse inconsistencies. In some situations, it may be more desirable for a user to clearly identify data that are in conflict before any decision or action is taken. Perhaps the inconsistencies have no bearing on the overall objective of the user, e.g. it is unlikely that an inconsistent description of the colour of the seats in an aircraft is relevant to the overall safety of the aircraft. Even in cases where corrective measures must be taken toward inconsistencies, it is unclear that a single action would suffice. In some cases, corrections must be performed gradually over time and there may be a need to provide a more quantitative way to monitor the progress of the repair.

1.4 Preservation and Measuring Inconsistent Information

Our main strategy for integrating the symbolic and numeric approaches, the logic as deduction and logic as representation views of knowledge representation, is to develop a very general methodology for comparing different inconsistency tolerant inference mechanisms *quantitatively*. This is a departure from standard methods for comparing inference mechanisms. Typically, comparisons between different inference mechanisms are drawn along the lines of

1. set theoretic relations (inclusion): we ask whether conclusions of one mechanism can be deduced by another mechanism for a given set of premises.

2. proof theory, i.e. axioms or derivable inference rules: we ask whether the axioms or inference rules of one mechanism can be derived by another mechanism.

3. computational complexity: we ask whether one mechanism is computationally more expensive (space and time) than another mechanism.

Our strategy here is rather different. We would like to provide a novel and useful way to compare inference mechanisms in terms of various quantitative properties or measurements that can be *preserved* from a given set of premises to conclusions. We envision that a key role of the analyst in our hybrid system is to provide such quantitative analyses of the input data. The key issue here is the notion of *preservation*. In the standard account of inference, the validity of an inference is defined in terms of the preservation of truth relative to the class of standard two-valued models – in a valid inference it is not possible for the premises to be true while the conclusion is false. Truths are transmitted from premises to conclusions in valid inferences. But this account is unhelpful for inconsistent premises since, according the standard two-valued semantics, inconsistent premises cannot be true together. So for inconsistent premises, there is simply no truth to be preserved. For information that is uncertain, we need a more pragmatic approach to the notion of preservation. The general idea is that, apart from the standard semantic or model theoretic properties, there may well be other metalogical properties that are transmitted from premises to conclusions in an inference. Presumably some of these properties would be of interest to a user of the system, depending of course on the user's overall objective. And on occasions, it may even be desirable to preserve these properties in an inference. So the role of the analyst here, at least in part, is to keep the user (and the prover) informed about the hidden structure and properties of the input data and perhaps to serve as an adviser for selecting the appropriate inference mechanism to preserve the properties of choice. Of course the requirement that an inference mechanism preserves more then truth in the standard models also implies that the mechanism, though sound, would not be complete with respect to standard valid inferences. Such a mechanism would in general preserve truth in the standard models but not all truth preserving inferences would be provable. The use of sound but incomplete as well as unsound but complete reasoning has already been investigated by Levesque in [119] from the point of view of computational complexity and Schaerf and Cadoli [162] from the point of view of approximate reasoning. But for us, the use of incomplete reasoning is an interesting *paraconsistent* approach to inconsistency. In general we agree that there is a need to adopt a weaker inference mechanism to perform deduction with inconsistent input data. But from a preservationist point of view adopting weaker inference mechanisms is not enough, we need to understand how different sound but incomplete inference mechanisms would preserve different properties of different inconsistent premises. In this respect we have at least two options within which to proceed.

1. we can identify sound but incomplete provers and investigate which additional properties they can preserve (or fail to preserve).

2. we can identify properties of inconsistent logical descriptions and investigate which prover can preserve (or fail to preserve) these properties.

These two options are really two aspects of the same picture. Viewed more abstractly (see figure (1.4)), the preservational approach takes deduction to be an operation, C, defined over a formal language Φ, i.e. $C : \wp(\Phi) \longrightarrow \wp(\Phi)$. Given a (partially or totally) ordered set $(S, \leq)$, a metalogical property is simply a function from $\wp(\Phi)$ to S, i.e. $f : \wp(\Phi) \longrightarrow S$. From the preservational point of view, the crucial question is: given an arbitrary $\Gamma \subseteq \Phi$, what is the relation between $f(C(\Gamma))$ and $f(\Gamma)$ in terms of the ordering $\leq$?

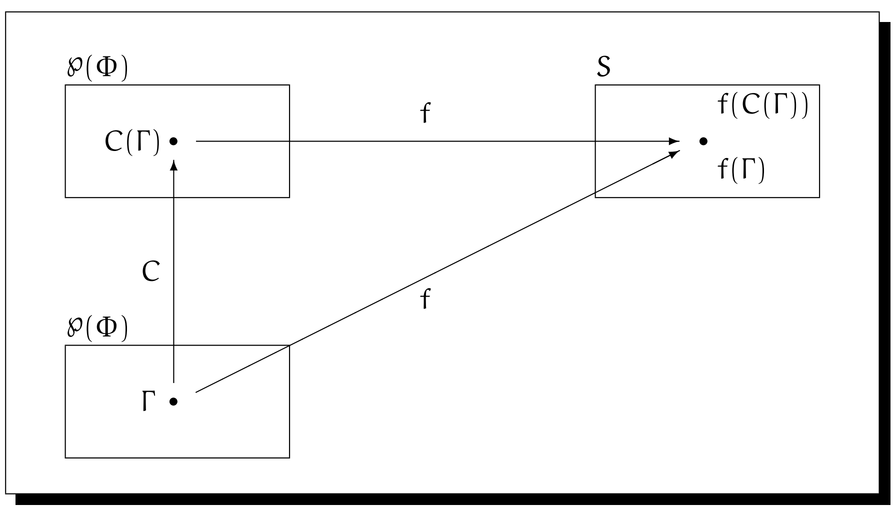

Figure 1.4: An abstract view of Preservation

It is easy to see that this abstract view of preservation captures the standard notion of the soundness and completeness of a logic. Let C be the closure under deduction of a logic L, and let members of S be collections of models defined according to a fixed semantics, take the ordering on S to be the usual inclusion ordering $\subseteq$. Let f be the function which assigns to each $\Gamma \subseteq \Phi$, the collection of models for Γ, i.e. every element of the collection is a model of Γ. Then to say that the logic L is sound and complete with respect to S is precisely to say that $f(C(\Gamma)) = f(\Gamma)$ for every $\Gamma \subseteq \Phi$. The soundness of L is captured by the inclusion $f(\Gamma) \subseteq f(C(\Gamma))$, whereas the completeness of L is captured by the inclusion, $f(C(\Gamma)) \subseteq f(\Gamma)$.

This abstract view of preservation is extremely minimal. In figure (1.4), we make no assumption about the nature of the set S or the function f. But different f and S would in fact give us different ways to partition $\wp(\Phi)$. For each $x \in S$, the set $f^{-1}[x] = \{\Gamma \subseteq \Phi : f(\Gamma) = x\}$ obviously forms an equivalence class. Similarly, this is true with respect to the deductive closure C. Since we are interested in both inference and analysis, we are interested in both deductive closures C and functions f that can distinguish between different inconsistent sets. Moreover, since we want to provide quantitative ways to distinguish inconsistent sets, our interest is in those functions f that range over different numerical sets S, e.g. S may simply be $\mathbb{N}$, $[0, 1]$ or even $[0, 1]^n$.

To take one particular example from probabilistic inference, the *uncertainty* of a proposition A is defined as by

$$U(A) = 1 - P(A) \tag{1.1}$$

In (1.1), $P(A)$ is the probability that A is true. It is a straightforward consequence of the Kolmogorov Axioms for probability that if B is deducible from A in classical propositional or first order logics, then $U(B) \leq U(A)$, i.e. $U(B) \not> U(A)$ given that $\leq$ is the usual total ordering on $[0, 1]$. In particular, this means that for any classical inference if we begin with a single premise with small uncertainty (high certainty), then any one of its conclusions can only have small uncertainty (high certainty). Conversely, the uncertainty of a conclusion is large only if the uncertainty of its (single) premise is large (see [3; 4] for more details). Accordingly, in the limiting case in which the uncertainty of a premise A is zero, $U(A) = 0$, then the uncertainty of its conclusion B must be zero, $U(B) = 0$. Not surprisingly, this just is the classical notion of deductive validity.

Clearly, the uncertainty measure U is one of many different properties that can be transmitted from premises to conclusions in an inference. But as we have suggested earlier, the preservation approach to inference is sensitive to the issue of developing theoretically meaningful ways to measure, compare and analyse inconsistent data. This is particularly important when we are using logic as a representation. Consider again a very simple example in which we are modeling negotiation by agents with conflicting objectives. Suppose we have 3 agents negotiating over 3 issues represented by a set of propositional variables $\{p_1, p_2, p_3\}$. In the first round of the negotiation the agents' positions are represented by

$$\Gamma_0 = \{p_1 \wedge p_2 \wedge p_3,\ \neg p_1 \wedge p_2 \wedge \neg p_3,\ p_1 \wedge \neg p_2 \wedge \neg p_3\}$$

In the second round the agents' positions are represented by

$$\Gamma_1 = \{p_1 \wedge p_2 \wedge p_3,\ \neg p_1 \wedge p_2 \wedge p_3,\ p_1 \wedge \neg p_2 \wedge p_3\}$$

Are there significant differences between Γ_0 and Γ_1? The answer is 'no' if we take C to be the closure under deduction of classical logic and f to be the standard mapping of sets of formulae to their classical models. As far as classical C and standard f are concerned, Γ_0 and Γ_1 belong to the same equivalence class. But if we are interested in monitoring the progress of the negotiation, then the answer is a definite 'yes'. Clearly in the second round of the negotiation the agents reach an agreement about p_3 even though there is no general agreement about p_1 or p_2. The take home message of our example is this: even at a very simple propositional level different inconsistent logical descriptions are endowed with very different combinatorial structures and properties. An inference mechanism that fails to differentiate between Γ_0 and Γ_1 also fails to recognize potentially important and useful information for the user. Thus in the design of an inference mechanism that can handle and tolerate inconsistencies, we need to pay attention to these underlying combinatorial structures and properties.

What kind of analysis must the analyst in our hybrid system provide in order to distinguish between Γ_0 and Γ_1? Since a main assumption in this thesis is that logical descriptions are information bearing, a natural starting point would seem to be the traditional information theory of Shannon – it provides a theoretical foundation for measuring the amount of information in a set of data. But as it stands, the standard information theory is inadequate since it typically treats inconsistent data as containing either maximum amount of information or no information at all (see [6] for instance). In this respect, we are interested in both extending standard information theory to cover inconsistent data and in finding alternative ways to measure and compare inconsistent data. These alternatives are more or less what we may call a divide and conquer method for handling inconsistency. The basic strategy is to divide an inconsistent set into (not necessary consistent) subsets, depending on how we make the cut, we get different ways to measure and classify sets.

1.5 Representation of Inconsistent Information

What makes one quantitative analysis more fruitful and meaningful then another? There is no fast and easy answer to this question. The answer to the question is largely dependent on the underlying objective of the user in possession of the logical description. But minimally, it would be desirable if the analysis can be reused across different representational formalisms. This would show that the properties so specified are invariant under different formalisms, and hence that they are not just incidental features of a particular formalism. One way in which we can demonstrate this is to actually apply the concepts and analyses developed for one formalism to another. In this respect, there are two natural ways to proceed. Since inconsistencies often occur within *intensional* linguistic contexts involving various form of modalities, a natural way to extend the expressive power of a propositional language to introduce modal operators for the corresponding doxastic, epistemic or deontic contexts.[1] This gives us a direct way to study inconsistency within a modal context. But doing so also raises an interesting issue about the choice of modal logics. The weakest modal logic adequate for Kripkean binary relational semantics is the logic K. Although in K we can provide a binary relational model for $\{\Box A,\ \Box\neg A\}$, the following rule is derivable in K and any of its extension:

$$[\Box\ ECQ] \qquad \frac{\Box A \qquad \Box\neg A}{\Box B}$$

So in any binary relational model adequate for the logic K, any world which verifies $\{\Box A, \Box\neg A\}$ would be a world which also verifies $\Box B$ for any B. For the purpose of modelling intensionalized inconsistencies then, the standard modal logics and their corresponding Kripkean semantics seem to be out of place. How can we avoid the rule $[\Box\ ECQ]$? One possible option is to adopt modal logics that are strictly weaker then K and to develop alternative semantics for these logics. In particular, we need a provide semantics for modeling $\{\Box A, \Box\neg A\}$ without also modelling $\Box B$ for any B. This is the approach taken by Fagin and Halpern [67], Jaspars [98], Massacci [127] and Rantala [148]. In this thesis, we'll be looking at weaker modal logics for the representation of intensionalised inconsistencies.

While extending propositional language with modalities is one direction which we may take to explore representational issues of inconsistency, another direction is diagrammatic reasoning systems where information is encoded not as a linear sequence of symbols but as two dimensional figures or diagrams. Indeed many forms of propo-

[1] See [75] for more details concerning the distinction between *extensional* and *intensional* contexts.

sitional reasoning can be represented as graphs or hypergraphs. So it is natural for us to explore these formalisms for representing inconsistent information.

1.6 Overview

In chapter two, we'll carry out the preservational approach to analyse various standard inference mechanisms based on reasoning from consistent subsets. We'll introduce various quantitative measurements for inconsistency and investigate the preservational properties of these mechanisms in light of these measurements.

In chapter three, we'll highlight the well known connection between reasoning from maximal consistent subsets and the standard default reasoning developed by Reiter. The implication is that the kind of preservational analyses offered in chapter three have direct counterparts in default reasoning. We'll also address a criticism offered by Belnap against reasoning based on maximal consistent subsets (and hence, indirectly against default reasoning). We'll point out that Belnap's own amendment does not in fact resolve the very difficulty he raised. We'll propose an amendment to Belnap's' amendment.

In chapter four, we'll look at the issue of preservation from the point of view of *uncertainties* that are transmitted from premises to conclusions. As we have already seen in the single premise case the uncertainty of a classical conclusion is always bound by the uncertainty of the premise. But for a set of premises this is no longer true for classical logic. Although the uncertainty of each premise in a set may be small, the uncertainty of a conclusion may turn out to be prohibitively high. In this chapter, we'll investigate uncertainty phenomena in light of inconsistency. This leads to some surprising results and conjectures.

In chapter five, we'll introduce the paraconsistent logic QC developed originally by Besnard and Hunter in [34; 91; 92]. We outline a particular strategy to use (half of) QC logic as an assistant to analyse inconsistent data. We'll also consider using this strategy to study over-constrained problems.

In chapter six, we'll investigate the use of modal logics for representing inconsistency. The family of modal logics presented here is a generalization of those developed by Jennings and Schotch in [100; 101; 164; 165]. Instead of treating modality as an unary operator, we take modality to be a multi-ary operator. The models developed here combine both relational semantics and neighbourhood semantics. A completeness proof is given utilising the technique developed by Brown and Apostoli [11; 9;

10].

In chapter seven, we'll look at a hypergraph representation of a covering theoretic measurement of inconsistency. We'll develop a general notion of n-satisfiability on hypergraphs and show that the compactness statement of n-satisfiability on hypergraphs is equivalent to BPI in ZF set theory. We give a syntactic characterization of n-satisfiability on hypergraphs in terms of a resolution style proof procedure. A general notion of consequence relations based on hypergraphs will also be introduced. We'll conclude with a discussion of a conjecture of Cowen relating BPI and complexity theory.

Finally in chapter eight, we present the conclusions and directions for future work. For completeness, we have also included several appendices at the end.

1.7 Notation

In subsequent discussions we'll assume the simplest logical language – propositional language – and examine various inconsistent tolerant formal mechanisms therein. We assume that Φ is a set of propositional formulae generated from propositional atoms or variables, $\{p_1, q_1, p_2, q_2, \ldots\}$, with the usual boolean connectives, $\neg, \wedge, \vee, \supset$. We use A, B, C, $\ldots$, to denote formulae, $\top$ for any tautology, $\bot$ for any contradiction, Γ, Σ, Δ, $\ldots$, to denote sets of formulae, and $\mathcal{A}$, $\mathcal{B}$, $\ldots$, to denote subsets of a set of formulae. From time to time we'll use 'Γ, Δ' and '$\Gamma \cup \Delta$' interchangeably especially if Δ is a singleton. We assume the equivalence between $A \supset B$ and $\neg A \vee B$.

There are many extensionally equivalent ways to characterize classical logic. We give the standard Hilbert style axiomatic definition here. By a *deduction* of A from a (possibly infinite) set Γ, we mean a finite sequence $\langle A_1, \ldots, A_n \rangle$ such that $A_n = A$ and for each $i \leq n$, A_i is either an axiom, a member of the set Γ, or obtained by modus ponens from two previous formulae A_j and A_k where $j < k < i$. We use $\vdash$ to denote the classical deducibility relation and $\mathbf{Cn}(\Gamma)$ to denote $\{A \in \Phi : \Gamma \vdash A\}$. A set of formulae Γ is inconsistent if $\Gamma \vdash \bot$, otherwise Γ is consistent. A theory $\mathcal{T}$ is a set of formulae that is closed under $\vdash$, i.e. $A \in \mathcal{T}$ iff $\mathcal{T} \vdash A$.

From time to time we'll also make use of the fact that **Cn** is a compact Tarskian closure operator over Φ, i.e. **Cn** has the following properties:

Inclusion $\Gamma \subseteq \mathbf{Cn}(\Gamma)$

Monotonicity $\Gamma \subseteq \Delta \implies \mathbf{Cn}(\Gamma) \subseteq \mathbf{Cn}(\Delta)$

Idempotence $\mathbf{Cn}(\mathbf{Cn}(\Gamma)) = \mathbf{Cn}(\Gamma)$

Compactness $\mathbf{Cn}(\Gamma) = \Phi \implies \mathbf{Cn}(\Gamma') = \Phi$ for some $\Gamma' \subseteq_{fin} \Gamma$

As usual we'll use the standard set theoretic abstraction notation $\{x | P(x)\}$ for the set of objects that has the property P. Notations of operations and functions defined on sets are given in the usual way. Where f and g are functions defined over the same sets, we use f^{-1} to denote the *inverse* of f, i.e. $f^{-1} = \{\langle y, x\rangle : \langle x, y\rangle \in f\}$. $f \circ g$ is the *composition* of g and f, i.e. for each x, $f \circ g(x) = g(f(x))$. $f \upharpoonright A$ is the *restriction* of f to A, i.e. $f \upharpoonright A = \{\langle x, f(x)\rangle : x \in A\}$. $f[A]$ is the *image* of A under f, i.e. $f[A] = \{f(x) : x \in A\}$. Note that inverse, composition, restriction and image need not apply only to functions, they can be defined for relations as well. We take an *injection* to be a $1-1$ function, a *surjection* to be an onto function, and a *bijection* to be an injective and surjective function. A *cartesian product* of A and B, denoted by $A \times B$, is defined by setting $A \times B = \{\langle x, y\rangle | x \in A \wedge y \in B\}$. A binary relation R defined on A is any subset of the cartesian product $A \times A$ (or A^2). For readability we write xRy or Rxy to denote $\langle x, y\rangle \in R$. A binary relation R defined on A is *reflexive* if for all $x \in A$, xRx. R is *symmetric* if for all $x, y \in A$, $xRy \rightarrow yRx$. R is *antisymmetric* if for all $x, y \in A$, $(xRy \wedge yRx) \rightarrow x = y$. R is *connected* if for all $x, y \in A$, $x \neq y \rightarrow (xRy \vee yRx)$. R is *transitive* if for all $x, y, z \in A$, $(xRy \wedge yRz) \rightarrow xRz$. A binary relation is an *equivalence relation* if it is reflexive, symmetric and transitive. A binary relation is a *partial ordering* if it is reflexive, antisymmetric and transitive. A binary relation is a *total ordering* or *linear ordering* if it is connected partial ordering relation. A *partial ordering set* or *poset* is a pair $\langle A, \leq\rangle$ where $\leq$ is a partial ordering defined on A. A poset $\langle A, \leq\rangle$ is *well founded* if every nonempty subset of A has a $\leq$-minimal element (equivalently there is no infinitely descending $\leq$-chain). A *total ordering set* or *toset* is a pair $\langle A, \leq\rangle$ where $\leq$ is a total ordering defined on A. A *well ordering set* or *woset* is a well founded poset.

Chapter 2

Paraconsistent Inference and Preservation

2.1 Introduction

Correct reasoning is usually characterised as patterns of inference which preserve truth. According to the standard view an inference is valid if it is impossible for its premises to be true but its conclusion false. While not incorrect, the standard view is unhelpful when we are confronted with inconsistent data. Since all inconsistent sets are unsatisfiable in the standard two-valued semantics, inferences licensed by classical logics become unprincipled in the presence of inconsistencies.

Many proposals and remedies are available to achieve inconsistency tolerant reasoning. They include both semantic and syntactic approaches:

1. introduce additional truth values to alter the semantics [12; 21; 37; 138]
2. introduce additional semantic parameters such as nonstandard possible worlds, setups or situations to evaluate formulae [68; 155; 157]
3. introduce labels or annotations into the object language, typically attached to formulae, to represent inconsistencies [41; 108; 124]

Undoubtedly, many semantic and syntactic innovations are involved in these approaches. In [99], Jennings *et al* have proposed a more pragmatic account of reasoning according to which the aim of logic is to provide a theory of reasoning which specifies the procedures for preserving important metalinguistic properties of premise sets. Accordingly, a practical reasoning system provides procedures by which a set of sentences having some metalinguistic properties can be unfailingly extended to a larger set with the same properties.

From this preservation-theoretic framework we can articulate two strategies for studying reasoning: the first is the identification of important metalinguistic properties of premises, and the second is the discovery of mechanisms that preserve these properties. Provisionally, no restriction is imposed on the kind of properties to be studied, except the properties in question must be strictly non-monotonic, i.e. not closed under supersets. Our main objective in this chapter is to carry out a program of research which takes the notion of preservation seriously and to give an analysis of various inconsistency tolerant reasoning strategies therein. Note that in stating our objective, we have implicitly endorsed a set theoretic presentation of premises and conclusions. But this assumption is inessential to the underlying methodological point. If premises and conclusions are modelled as different abstract data types, e.g. as multisets or lists, we can rephrase all our definitions accordingly. In any case, we are interested in inconsistency-tolerant inferences, whatever ways premises and conclusions are presented, and their preservational properties.

2.2 Paraconsistent Inferences

One common approach to handling inconsistencies resulting from information fusion from multiple sources is to fragment an inconsistent set into maximal consistent subsets and then extract conclusions by applying classical inference to these subsets. This approach was first introduced by Rescher and Manor [152; 153; 154; 156] and more recently extended by Benferhat *et al* [25; 27; 26; 28; 29; 32; 33; 31; 30].

In this section we present similar but slightly more general inference mechanisms to extract conclusions from an inconsistent set. As we shall see in the next chapter, the inference strategies presented here are expressively equivalent to reasoning from maximal consistent subsets. In our framework, an inference is a ternary relation between a premiss set Γ, a consistent *constraint* set Σ and a conclusion A.

Definition 2.2.1
Let Γ be a premiss set and Σ be an arbitrary but fixed consistent set which we call a constraint set on Γ. Then a subset $\mathcal{A}$ of Γ is Σ-inconsistent iff $\mathcal{A} \cup \Sigma$ is inconsistent, else $\mathcal{A}$ is Σ-consistent. A maximal Σ-consistent subset of Γ is a subset of Γ which has no proper Σ-consistent extension.

The main motivation behind definition (2.2.1) is that we need to be able to distinguish between different types of information in certain reasoning tasks. For instance, some information may have higher priority than others or it may provide us with specific knowledge of a domain. The role of Σ in definition (2.2.1) is to rule out data in Γ that is *bad* relative to Σ. Intuitively, we may think of Σ as a set of secured or prioritized data, or even just a set of background beliefs of an agent at a given time. But more concretely, it is similar to the idea of integrity constraints in database theory where the known relationships between various data elements are specified. To illustrate, consider the following example:

Example 2.2.1
The following information about a particular individual is obtained through a questionnaire:

```
marital status = married
        age = 1
```

In countries in which the legal age for marriage is 18, the above information is inconsistent.[1] In such cases, we need to check our data against the following integrity constraint:

$$\texttt{marital status = married} \supset \texttt{age > 17}$$

So the idea to view an inference as a ternary relation makes sense both theoretically and practically from the point of view of information processing. In many knowledge base and knowledge representation systems there are additional restrictions imposed on the *knowledge base language* used to express information available in Γ, the integrity constraint Σ and the *query language* for expressing the inferred conclusion A. In a relational database for instance, Γ is a set of *negation free* positive facts and integrity constraints in Σ are negated closed formulae. In a deductive database, Γ can contain either positive facts or rules whose heads are atoms and bodies are literals. These restrictions not only provide a greater degree of control over inference in terms of what can be derived from what, but in some cases they are indispensable to reducing the complexity of inferences (see Wagner [179] for more discussion).

For our purpose, we'll assume that Σ is an arbitrary but fixed constraint set. The set of all maximal Σ-consistent subsets of Γ is denoted by $M_\Sigma(\Gamma)$. Given a Σ-inconsistent premiss set Γ, an element $A \in \Gamma$ is a Σ-witness if $\{A\}$ is Σ-consistent, otherwise A is a Σ-villain. We define the safe part of Γ as, $S_\Sigma(\Gamma) = \bigcap M_\Sigma(\Gamma)$. We say that a subset

[1] In different counteries the integrity contraint may be different. As far as we know, no country currently permit legal marriage of children of age one

$\mathcal{A} \subset \Gamma$ is large iff $\mathcal{A} \in M_\Sigma(\Gamma)$ and for each $\mathcal{B} \in M_\Sigma(\Gamma)$, $|\mathcal{B}| \leq |\mathcal{A}|$. We use $L_\Sigma(\Gamma)$ to denote the set of all large subsets of Γ. A subset $\mathcal{A}$ of Γ is a minimally Σ-inconsistent subset if it is Σ-inconsistent and no proper subset of $\mathcal{A}$ is Σ-inconsistent. The set of all minimally Σ-inconsistent subset of Γ is denoted by $MI_\Sigma(\Gamma)$. The Σ-inconsistent part of Γ is defined by:

$$In_\Sigma(\Gamma) = \bigcup MI_\Sigma(\Gamma)$$

A set $\mathcal{H}$ is a *hitting set* of a collection of sets, $\mathcal{C} = \{\mathcal{S}_i : i \in I\}$, if for every $i \in I$, $\mathcal{S}_i \cap \mathcal{H} \neq \emptyset$. $\mathcal{H}$ is a *minimal hitting set* if none of its proper subsets are hitting sets of $\mathcal{C}$. Given the notion of minimal hitting set, $M_\Sigma(\Gamma)$ and $MI_\Sigma(\Gamma)$ are interdefinable.

Proposition 2.2.1
Let Σ be an arbitrary but fixed constraint set and Γ be any premise set.

1. $\mathcal{A} \in M_\Sigma(\Gamma) \iff \Gamma \setminus \mathcal{A}$ *is a minimal hitting set of* $MI_\Sigma(\Gamma)$.
2. $\mathcal{B} \in MI_\Sigma(\Gamma) \iff \mathcal{B}$ *is a minimal hitting set of* $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$.
3. $|M_\Sigma(\Gamma)| \geq \max\{|\mathcal{B}| : \mathcal{B} \in MI_\Sigma(\Gamma)\}$, *moreover equality holds if* $|MI_\Sigma(\Gamma)| = 1$.
4. $|MI_\Sigma(\Gamma)| \geq \max\{|\Gamma \setminus \mathcal{A}| : \mathcal{A} \in M_\Sigma(\Gamma)\}$, *moreover equality holds if* $|M_\Sigma(\Gamma)| = 1$.

Proof:
(1. $\Rightarrow$) Suppose $\mathcal{A} \in M_\Sigma(\Gamma)$ but $\Gamma \setminus \mathcal{A}$ is not a hitting set of $MI_\Sigma(\Gamma)$. Then for some $\mathcal{B} \in MI_\Sigma(\Gamma)$, $(\Gamma \setminus \mathcal{A}) \cap \mathcal{B} = \emptyset$. This implies that $\mathcal{B} \subseteq \mathcal{A}$ which is impossible given the Σ-consistency of $\mathcal{A}$. Hence, $\Gamma \setminus \mathcal{A}$ must be a hitting set. Suppose that $\Gamma \setminus \mathcal{A}$ is not a minimal hitting set of $MI_\Sigma(\Gamma)$, then there must be a proper superset $\mathcal{A}' \supset \mathcal{A}$ such that $\Gamma \setminus \mathcal{A}'$ is a hitting set of $MI_\Sigma(\Gamma)$. Since $\Gamma \setminus \mathcal{A}'$ and $\mathcal{A}'$ are disjoint, $\mathcal{A}'$ cannot contain any $\mathcal{B} \in MI_\Sigma(\Gamma)$. This implies that $\mathcal{A}'$ is Σ-consistent which contradicts the maximality of $\mathcal{A}$.

(1. $\Leftarrow$) Let $\mathcal{A}$ be an arbitrary but fixed subset of Γ such that $\Gamma \setminus \mathcal{A}$ is a minimal hitting set of $MI_\Sigma(\Gamma)$. Then by the disjointness of Γ and $\Gamma \setminus \mathcal{A}$ for no $\mathcal{B} \in MI_\Sigma(\Gamma)$ do we have $\mathcal{B} \subseteq \mathcal{A}$. Hence $\mathcal{A}$ is Σ-consistent. Since every proper subset of $\Gamma \setminus \mathcal{A}$ is not a hitting set of $MI_\Sigma(\Gamma)$, every proper superset of $\mathcal{A}$ must contain some $\mathcal{B} \in MI_\Sigma(\Gamma)$. Hence $\mathcal{A}$ must be maximally Σ-consistent.

(2. $\Rightarrow$) Consider an arbitrary $\mathcal{B} \in MI_\Sigma(\Gamma)$. From (1) above, for every $\mathcal{A} \in M_\Sigma(\Gamma)$, $\Gamma \setminus \mathcal{A}$ is a minimal hitting set of $MI_\Sigma(\Gamma)$. So every $\Gamma \setminus \mathcal{A}$ must intersect $\mathcal{B}$. Hence $\mathcal{B}$ is a hitting set of $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$. Suppose that $\mathcal{B}$ is not minimal. Then there exists a proper

subset $\mathcal{B}' \subset \mathcal{B}$ such that $\mathcal{B}'$ is a hitting set of $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$. Since $\Gamma \setminus \mathcal{A}$ and $\mathcal{A}$ are disjoint for every $\mathcal{A} \in M_\Sigma(\Gamma)$, for no $\mathcal{A} \in M_\Sigma(\Gamma)$ do we have $\mathcal{B}' \subseteq \mathcal{A}$. But this is impossible since $\mathcal{B} \in MI_\Sigma(\Gamma)$, $\mathcal{B}'$ must be Σ-consistent. Hence, $\mathcal{B}$ must be a minimal hitting set.

(2. $\Leftarrow$) Let $\mathcal{B}$ be a minimal hitting set of $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$. Suppose that $\mathcal{B}$ is Σ-consistent. Then for some $\mathcal{A}_0 \in M_\Sigma(\Gamma)$, $\mathcal{B} \subseteq \mathcal{A}_0$. This implies that $\mathcal{B} \cap (\Gamma \setminus \mathcal{A}_0) = \emptyset$ which is impossible since $\mathcal{B}$ is a hitting set of $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$. Hence, $\mathcal{B}$ must be Σ-inconsistent. Suppose that $\mathcal{B}$ is not minimally Σ-inconsistent. Then there must be a proper subset $\mathcal{B}' \subset \mathcal{B}$ which is minimally Σ-inconsistent. From (2. $\Rightarrow$) above, $\mathcal{B}'$ is a minimal hitting set of $\{\Gamma \setminus \mathcal{A} : \mathcal{A} \in M_\Sigma(\Gamma)\}$. This contradicts the assumption that $\mathcal{B}$ is a minimal hitting set.

(3) If $\mathcal{B}$ is the largest minimal Σ-inconsistent subset of Γ, then for each $B \in \mathcal{B}$, $\mathcal{B} \setminus \{B\}$ is Σ-consistent. There are exactly $|\mathcal{B}|$ many such sets. Hence there are at least $|\mathcal{B}|$ many maximal Σ-consistent subsets of Γ. Clearly $|M_\Sigma(\Gamma)| = |\mathcal{B}|$ if $\mathcal{B}$ is the only minimal Σ-inconsistent subset.

(4) From (1) if $\mathcal{A} \in M_\Sigma(\Gamma)$, then $\Gamma \setminus \mathcal{A}$ is a minimal hitting set of $MI_\Sigma(\Gamma)$. So if $\Gamma \setminus \mathcal{A}$ is the largest such set, then there are at least $|\Gamma \setminus \mathcal{A}|$ many minimal Σ-inconsistent subsets of Γ since $(\Gamma \setminus \mathcal{A}) \cap \mathcal{B} \neq \emptyset$ for each distinct $\mathcal{B} \in MI_\Sigma(\Gamma)$. Clearly in the event that $\mathcal{A}$ is the only maximal Σ-consistent subset of Γ, we have $|MI_\Sigma(\Gamma)| = |\Gamma \setminus \mathcal{A}|$. ∎

We now define the following notions of consequence in terms of $M_\Sigma(\Gamma)$. Given proposition (2.2.1), each of the following consequences can be defined in terms of $MI_\Sigma(\Gamma)$ as well.

Definition 2.2.2

Σ-universal-consequence $A \in C_{U\Sigma}(\Gamma)$ *iff for each* $\mathcal{A} \in M_\Sigma(\Gamma)$, $\mathcal{A} \vdash A$

Σ-existential-consequence $A \in C_{E\Sigma}(\Gamma)$ *iff for some* $\mathcal{A} \in M_\Sigma(\Gamma)$, $\mathcal{A} \vdash A$

Σ-argued-consequence $A \in C_{A\Sigma}(\Gamma)$ *iff there exists some* $\mathcal{A}_i \in M_\Sigma(\Gamma)$ *with* $\mathcal{A}_i \vdash A$ *and for every* $\mathcal{A}_j \in M_\Sigma(\Gamma)$, $\mathcal{A}_j \nvdash \neg A$.

Σ-safe-consequence $A \in C_{S\Sigma}(\Gamma)$ *iff* $S_\Sigma(\Gamma) \vdash A$

Σ-large-consequence $A \in C_{L\Sigma}(\Gamma)$ *iff* $\mathcal{A} \vdash A$ *for each* $\mathcal{A} \in L_\Sigma(\Gamma)$.

As usual for $x \in \{S\Sigma,\ U\Sigma,\ A\Sigma,\ L\Sigma,\ E\Sigma\}$, we can define the corresponding inference relation $\vdash_x$ by setting $\Gamma \vdash_x A$ iff $A \in C_x(\Gamma)$. The following proposition is an easy consequence of our definitions.

Proposition 2.2.2

Let Σ be an arbitrary but fixed constraint set on Γ where $M_\Sigma(\Gamma) \neq \emptyset$,

1. $S_\Sigma(\Gamma) = \Gamma \setminus In_\Sigma(\Gamma)$
2. $C_{S\Sigma}(\Gamma) \subseteq C_{U\Sigma}(\Gamma) \subseteq C_{A\Sigma}(\Gamma) \subseteq C_{E\Sigma}(\Gamma)$
3. $C_{S\Sigma}(\Gamma) \subseteq C_{U\Sigma}(\Gamma) \subseteq C_{L\Sigma}(\Gamma) \subseteq C_{E\Sigma}(\Gamma)$

Proof:

(1) $S_\Sigma(\Gamma) \subseteq \Gamma \setminus In_\Sigma(\Gamma)$: Let $A \in S_\Sigma(\Gamma)$. Then $A \in \bigcap M_\Sigma(\Gamma)$. Suppose to the contrary that $A \in In_\Sigma(\Gamma)$. Then there must be some $\mathcal{B}_0 \in MI_\Sigma(\Gamma)$ such that $A \in \mathcal{B}_0$. But by the minimality, $\mathcal{B}_0 \setminus \{A\}$ must be Σ-consistent. So there must be some $\mathcal{A} \in M_\Sigma(\Gamma)$ such that $(\mathcal{B}_0 \setminus \{A\}) \subseteq \mathcal{A}$. But by the initial assumption $A \in \bigcap M_\Sigma(\Gamma)$ and so $A \in \mathcal{A}$. But then $(\mathcal{B}_0 \setminus \{A\}) \subseteq \mathcal{A}$ and $\{A\} \subseteq \mathcal{A}$. Hence $\mathcal{B}_0 \subseteq \mathcal{A}$ which contradicts the assumption that $\mathcal{B}_0 \in MI_\Sigma(\Gamma)$. Hence we must reject the assumption that $A \in In_\Sigma(\Gamma)$.
$\Gamma \setminus In_\Sigma(\Gamma) \subseteq S_\Sigma(\Gamma)$: Let $A \in \Gamma$ but $A \notin \bigcap M_\Sigma(\Gamma)$. Then there must be some $\mathcal{A} \in M_\Sigma(\Gamma)$ such that $\mathcal{A} \cup \{A\}$ is Σ-inconsistent, for otherwise $\mathcal{A} \cup \{A\}$ is Σ-consistent for every $\mathcal{A} \in M_\Sigma(\Gamma)$ and thus $A \in \bigcap M_\Sigma(\Gamma)$ contradicting the initial assumption. Hence there must be some $\mathcal{B}_0 \in MI_\Sigma(\Gamma)$ such that $\mathcal{B}_0 \subseteq \mathcal{A} \cup \{A\}$ and $A \in \mathcal{B}_0$. Hence $A \notin \Gamma \setminus In_\Sigma(\Gamma)$ as required.

(2) $C_{S\Sigma}(\Gamma) \subseteq C_{U\Sigma}(\Gamma)$: Clearly the containment holds since

$$\mathbf{Cn}(\bigcap M_\Sigma(\Gamma)) \subseteq \bigcap_{\mathcal{A} \in M_\Sigma(\Gamma)} \mathbf{Cn}(\mathcal{A})$$

$C_{U\Sigma}(\Gamma) \subseteq C_{A\Sigma}(\Gamma)$: we assume that $A \notin C_{A\Sigma}(\Gamma)$ and show that $A \notin C_{U\Sigma}(\Gamma)$. By the definition of $C_{A\Sigma}$, $A \notin C_{A\Sigma}(\Gamma)$ implies that $A \in C_{E\Sigma}(\Gamma)$ and $\neg A \in C_{E\Sigma}(\Gamma)$. Let $\mathcal{A}_0 \in M_\Sigma(\Gamma)$ be such that $\neg A \in \mathbf{Cn}(\mathcal{A}_0)$. Towards a contradiction we assume that $A \in \bigcap\{\mathbf{Cn}(\mathcal{A}) : \mathcal{A} \in M_\Sigma(\Gamma)\}$. Thus in particular $A \in \mathbf{Cn}(\mathcal{A}_0)$. But this is impossible since $\mathcal{A}_0$ is Σ-consistent and thus $\mathcal{A}_0$ must be consistent. This contradicts our previous claims.
$C_{A\Sigma}(\Gamma) \subseteq C_{E\Sigma}(\Gamma)$: the containment follows directly from the definitions.

(3) $C_{U\Sigma}(\Gamma) \subseteq C_{L\Sigma}(\Gamma)$: since $L_\Sigma(\Gamma) \subseteq M_\Sigma(\Gamma)$,

$$\bigcap_{\mathcal{A} \in M_\Sigma(\Gamma)} \mathbf{Cn}(\mathcal{A}) \subseteq \bigcap_{\mathcal{A} \in L_\Sigma(\Gamma)} \mathbf{Cn}(\mathcal{A})$$

Hence $C_{U\Sigma}(\Gamma) \subseteq C_{L\Sigma}(\Gamma)$ as required. ∎

We note that for particular Σ and Γ, $C_{A\Sigma}(\Gamma)$ and $C_{L\Sigma}(\Gamma)$ may be incomparable, i.e. $C_{A\Sigma}(\Gamma) \not\subseteq C_{L\Sigma}(\Gamma)$ and $C_{L\Sigma}(\Gamma) \not\subseteq C_{A\Sigma}(\Gamma)$. To take an example $\Sigma = \emptyset$ and $\Gamma = \{p, \neg p, \neg p \vee q\}$. Clearly $q \in C_{A\Sigma}(\Gamma)$ but $q \notin C_{L\Sigma}(\Gamma)$.

The relative (set inclusion) ordering of Σ-consequences is summarised in figure (2.1). Downward arrows indicate proper set inclusions.

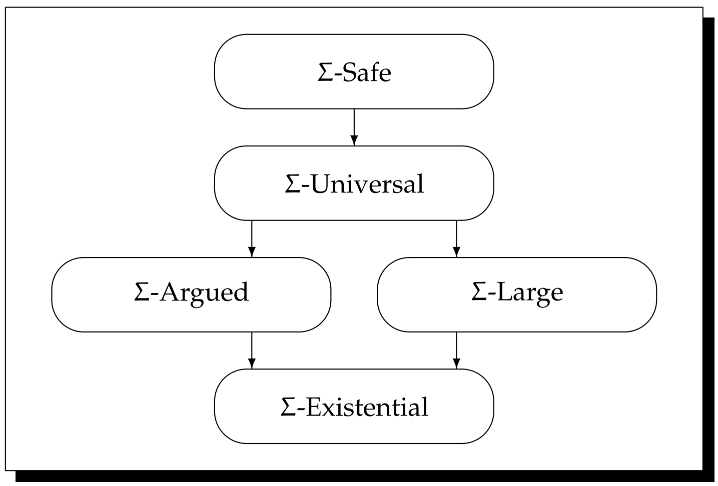

Figure 2.1: Inclusion ordering of Σ-consequences

From an inferential perspective, we can view $C_{S\Sigma}(\Gamma)$ and $C_{E\Sigma}(\Gamma)$ as reasoning strategies along a continuum. On the one hand, the Σ-safe consequence can be characterized as a species of skeptical inference since it regards any conflicting data as suspect and thus allows a reasoner to draw conclusions only from the safe part of a premiss set. Where the safe part of a premiss set is empty, $C_{S\Sigma}(\Gamma)$ contains only classical theorems. Σ-existential consequence, on the other hand, can be characterized as a species of liberal inference since it allows a reasoner to draw conclusions from any Σ-witness or cluster of Σ-witnesses of a set. So in the presence of both A and $\neg A$ in a set Γ, where A and $\neg A$ are both Σ-witnesses of Γ, $C_{E\Sigma}(\Gamma)$ contains both A and $\neg A$ individually (but not $A \wedge \neg A$). With respect to $C_{U\Sigma}(\Gamma)$, it is more liberal then $C_{S\Sigma}(\Gamma)$ but still remains cautious overall by accepting only the intersection of the classical consequences of all maximal Σ-consistent subsets of a premiss set. As for $C_{A\Sigma}(\Gamma)$, the main idea is to accept only those conclusions which we have direct arguments in their favour and no direct arguments for rejecting them. The notion of an argued-consequence forms the basis of the notion of an *argument system* which has been studied extensively in recent years ([36; 136; 159]). Argument systems have a game-theoretic flavour which makes them particularly suitable for modelling a variety of multi-agent systems ([160; 161; 177]).

One of the most peculiar features of Σ-argued consequence is captured in proposition (2.2.3). It highlights the fact that although an argued consequence has no direct refutation, each may still be rebutted when evidence is pooled together in some sense.

Proposition 2.2.3

1. *For each* $A \in C_{A\Sigma}(\Gamma)$*,* A *is* Σ*-consistent but* $C_{A\Sigma}(\Gamma)$ *is not* Σ*-consistent in general.*

2. *If* $\Sigma = \emptyset$*, then* $C_{A\Sigma}(\Gamma)$ *is pairwise* Σ*-consistent.*

Proof:
(1) If $A \in C_{A\Sigma}(\Gamma)$ is Σ-inconsistent, then $\Sigma \vdash \neg A$ and for some $\mathcal{A} \in M_\Sigma(\Gamma)$ $\mathcal{A} \vdash A$. Hence $\Sigma \cup \mathcal{A} \vdash \bot$, a contradiction.

To see that $C_{A\Sigma}(\Gamma)$ is not Σ-consistent in general, the following example suffices:

$$\begin{aligned} \Sigma &= \{r \wedge s\} \\ \Gamma &= \{p \wedge q,\ \neg p \wedge q\} \\ \Delta &= \{p \vee \neg r, \neg p \vee \neg s\} \end{aligned}$$

We can easily verify that every member of Δ is a Σ-argued consequence of Γ but Δ is not Σ-consistent.
(2) To see that $C_{A\emptyset}(\Gamma)$ is pairwise ∅-consistent, it suffices to observe that if $A, B \in C_{E\emptyset}(\Gamma)$ are such that $\{A, B\}$ is ∅-inconsistent, then there must be distinct $\mathcal{A}, \mathcal{B} \in M_\emptyset(\Gamma)$ such that $\mathcal{A} \vdash A$ and $\mathcal{B} \vdash B$ but $\mathcal{A} \vdash \neg B$ and $\mathcal{B} \vdash \neg A$. Hence $A, B \notin C_{A\Sigma}(\Gamma)$, a contradiction. ∎

Finally for $C_{L\Sigma}(\Gamma)$, the main idea is to accept only the intersection of the classical consequences of all Σ-large subsets of Γ, i.e.

$$C_{L\Sigma}(\Gamma) = \bigcap_{\mathcal{A} \in L_\Sigma(\Gamma)} \mathbf{Cn}(\mathcal{A})$$

Note that the notion of largeness naturally induces a total ordering on $M_\Sigma(\Gamma)$. This gives us a method to combine different strategies to obtain different consequences. For instance, we can define a new consequence $C_{SL\Sigma}$ by setting

$$C_{SL\Sigma}(\Gamma) = \mathbf{Cn}(\bigcap L(\Gamma))$$

Similarly, $C_{EL\Sigma}$ and $C_{AL\Sigma}$ can be defined accordingly. More generally, given a total or partial ordering $\leq$ on $M_\Sigma(\Gamma)$ we can apply any one of $C_{S\Sigma}$, $C_{U\Sigma}$, $C_{A\Sigma}$, and $C_{E\Sigma}$ to the $\leq$-maximal (or the $\leq$-minimal) elements of $M_\Sigma(\Gamma)$ to obtain a variety of consequences. The introduction of orderings forms the basis for a variety of *preferential* systems and semantics. Priest's LPm for instance is a paraconsistent logic whose consequence relation is defined in terms of the selection of LP models that are minimal with respect to the usual $\subseteq$ ordering (see [138] for more details).

We note that by setting $\Sigma = \emptyset$, we recover the paraconsistent consequences defined in [29]. We also note that in our definitions Σ only provides side constraints on premises. We can in fact allow Σ to be used directly to derive conclusions. We have the following stronger notions of consequence:

Definition 2.2.3

Σ-universal-consequence* $A \in C^*_{U\Sigma}(\Gamma)$ *iff for each* $\mathcal{A} \in M_\Sigma(\Gamma)$, $\mathcal{A} \cup \Sigma \vdash A$

Σ-existential-consequence* $A \in C^*_{E\Sigma}(\Gamma)$ *iff for some* $\mathcal{A} \in M_\Sigma(\Gamma)$, $\mathcal{A} \cup \Sigma \vdash A$

Σ-argued-consequence* $A \in C^*_{A\Sigma}(\Gamma)$ *iff for every* $\mathcal{A}_j \in M_\Sigma(\Gamma)$, $\mathcal{A}_j \cup \Sigma \nvdash \neg A$ *and there exists some* $\mathcal{A}_i \in M_\Sigma(\Gamma)$ *such that* $\mathcal{A}_i \cup \Sigma \vdash A$.

Σ-safe-consequence* $A \in C^*_{S\Sigma}(\Gamma)$ *iff* $S_\Sigma(\Gamma) \cup \Sigma \vdash A$

Σ-large-consequence* $A \in C^*_{L\Sigma}(\Gamma)$ *iff* $\mathcal{A} \cup \Sigma \vdash A$ *for each* $\mathcal{A} \in L_\Sigma(\Gamma)$.

It is easy to verify the following proposition:

Proposition 2.2.4

1. *For* $x \in \{U\Sigma,\ E\Sigma,\ S\Sigma,\ L\Sigma\}$, $C_x(\Gamma) \subseteq C^*_x(\Gamma)$.
2. *For some* Γ *and* Σ, $C_{A\Sigma}(\Gamma)$ *and* $C^*_{A\Sigma}(\Gamma)$ *are incomparable.*
3. $C^*_{A\Sigma}(\Gamma)$ *is pairwise* Σ*-consistent, but not* Σ*-consistent in general.*

Proof:

(1) Follows from the fact that **Cn** is monotonic.

(2) The following example suffices:

$$\Sigma = \{r\} \qquad\qquad \Gamma = \{p \wedge q,\ \neg p \wedge q,\ \neg q\}$$

It is straightforward to verify that

$$(\neg q \vee s) \wedge r \in C^*_{A\Sigma}(\Gamma) \quad \neg q \vee \neg r \notin C^*_{A\Sigma}(\Gamma)$$
$$(\neg q \vee s) \wedge r \notin C_{A\Sigma}(\Gamma) \quad \neg q \vee \neg r \in C_{A\Sigma}(\Gamma)$$

(3) To see that $C^*_{A\Sigma}(\Gamma)$ is pairwise Σ-consistent, it suffices to observe that if $A, B \in C^*_{E\Sigma}(\Gamma)$ are such that $\{A, B\}$ is Σ-inconsistent, then there must be distinct $\mathcal{A}, \mathcal{B} \in M_\Sigma(\Gamma)$ such that $\Sigma \cup \mathcal{A} \vdash A$ and $\Sigma \cup \mathcal{B} \vdash B$ but $\Sigma \cup \mathcal{A} \vdash \neg B$ and $\Sigma \cup \mathcal{B} \vdash \neg A$. Hence $A, B \notin C_{A\Sigma}(\Gamma)$, a contradiction.

To see that $C^*_{A\Sigma}(\Gamma)$ is not Σ-consistent in general, the following example suffices:

$$\begin{aligned}
\Sigma &= \{t\} \\
\Gamma &= \{p \wedge r,\ \neg p \wedge r, \neg r\} \\
\Delta &= \{(p \wedge r) \vee (q \wedge s),\ (\neg q \wedge r) \vee (\neg q \wedge s),\ (\neg r \vee \neg s)\}
\end{aligned}$$

We can verify that every member of Δ is a Σ-argued consequence* of Γ but Δ is inconsistent and hence also Σ-inconsistent. Note that our example also shows that the structural rules known as monotonicity and transitivity fail for both $C_{A\Sigma}$ and $C^*_{A\Sigma}$. ∎

Moreover, the relative set inclusion ordering of the Σ-consequences* is analogous to figure (2.1).

2.3 Some Structural Properties

Although all our inference relations are defined in terms of the classical $\vdash$, strictly speaking they are not consequence relations in the Tarski-Scott sense. We follow the terminology of Kraus, Lehmann and Magidor in [116] and list some structural properties commonly used for comparing nonmonotonic systems in table (2.1). We'll use $\mid\!\sim$ to denote an arbitrary inference relation. Note that in stating these structural properties no assumption is made about the underlying syntax of the language. The intuitive contents of these structural rules are fairly straightforward. Reflexivity says that any member of a set of assumptions is deducible. Monotonicity says that previously deduced conclusions are deducible from any enlarged set of assumptions. Transitivity says that once a lemma is generated, 'cut and paste' of deductions is possible to generate new deductions. Truth says that all tautologies are deducible. Consistency says that all classically consistent sets remain consistent. Left logical equivalence says that classically equivalent assumptions can be interchanged in deductions. Cautious monotonicity says that assumptions can be safely accumulated if they are each deducible. Right weakening says that classical consequences of deducible conclusions are also deducible. Finally supraclassicality says that deducibility is an extension of classical inference.

$$\frac{A \in \Gamma}{\Gamma \mathrel{|\!\sim} A}\text{ [Reflexivity]} \qquad \frac{A \dashv\vdash B \qquad \Gamma, B \mathrel{|\!\sim} C}{\Gamma, A \mathrel{|\!\sim} C}\text{ [Left Logical Equivalence]}$$

$$\frac{\Gamma \mathrel{|\!\sim} A}{\Gamma, \Gamma' \mathrel{|\!\sim} A}\text{ [Monotonicity]} \qquad \frac{\Gamma \mathrel{|\!\sim} A \qquad \Gamma \mathrel{|\!\sim} B}{\Gamma, A \mathrel{|\!\sim} B}\text{ [Cautious Monotonicity]}$$

$$\frac{\Gamma, A \mathrel{|\!\sim} B \qquad \Gamma \mathrel{|\!\sim} A}{\Gamma \mathrel{|\!\sim} B}\text{ [Transitivity]} \qquad \frac{\Gamma \mathrel{|\!\sim} A \qquad A \vdash B}{\Gamma \mathrel{|\!\sim} B}\text{ [Right Weakening]}$$

$$\frac{\vdash A}{\Gamma \mathrel{|\!\sim} A}\text{ [Truth]} \qquad \frac{\Gamma \vdash A}{\Gamma \mathrel{|\!\sim} A}\text{ [Supraclassicality]}$$

$$\frac{\Gamma \mathrel{|\!\sim} \bot}{\Gamma \vdash \bot}\text{ [Consistency]}$$

Table 2.1: Some structural properties of inference.

Relative to a fixed consistent constraint set Σ and a premise set Γ with $M_\Sigma(\Gamma) \neq \emptyset$, we can summarise the structural properties of our inference relations with table (2.2). The properties of the corresponding $*$ versions are completely similar. We use '$(*)$' to denote both versions of an inference relation, '$+$' and '$-$' to indicate that a structural rule holds or fails to hold respectively. The proof is routine calculation; we leave it to the reader.

The failure of transitivity in the case of $\vdash_{E\Sigma}$ and $\vdash_{A\Sigma}$ highlights an important conceptual distinction between these two inference strategies on the one hand and the remaining strategies on the other. Implicit in cases of $\vdash_{S\Sigma}$, $\vdash_{U\Sigma}$ and $\vdash_{L\Sigma}$ are the selection of a single Σ-consistent set of assumptions that are either implicitly or explicitly represented by Γ. We can think of these assumptions as the set of *available assumptions*. Once the selection is completed, all permissible deductions are restricted to the use of these available assumptions. In other words, there is a single set of Σ-consistent available assumptions fixed for all permissible deductions in these cases. Cutting and pasting of permissible deductions are thus also permissible since the ag-

	$\vdash^{(*)}_{S\Sigma}$	$\vdash^{(*)}_{U\Sigma}$	$\vdash^{(*)}_{A\Sigma}$	$\vdash^{(*)}_{L\Sigma}$	$\vdash^{(*)}_{E\Sigma}$
Reflexivity	−	−	−	−	−
Monotonicity	−	−	−	−	+
Transitivity	+	+	−	+	−
Left Logical Equivalence	+	+	+	+	+
Right Weakening	+	+	+	+	+
Truth	+	+	+	+	+
Consistency	+	+	+	+	+
Cautious Monotonicity	+	+	−	+	+
Supraclassicality	−	−	−	−	−

Table 2.2: Some structural properties of Σ-consequences.

gregate of the assumptions used are always a subset of the set of Σ-consistent available assumptions. So transitivity holds for $\vdash_{S\Sigma}$, $\vdash_{U\Sigma}$ and $\vdash_{L\Sigma}$. This is however not the case for $\vdash_{E\Sigma}$ and $\vdash_{A\Sigma}$. The set of available assumptions in these two cases are not Σ-consistent even though assumptions used in any given permissible deduction form a Σ-consistent subset. Cutting and pasting of permissible deductions may result in impermissible deduction since the aggregate of the assumptions used may turn out to be Σ-inconsistent. Hence transitivity fails for both $\vdash_{E\Sigma}$ and $\vdash_{A\Sigma}$.

2.4 Properties of Sets

In this section, we introduce two different properties of inconsistent sets. The first allows us to measure the relative *level of incoherence* of a premise set. The second provides a measurement of the relative *quantity of empirical information* of a premise set.

2.4.1 Level of Incoherence

Some inconsistent sets are clearly more unstable or incoherent then others. Consider, for instance,

Example 2.4.1
$\Gamma = \{p \wedge q,\ \neg p \wedge q,\ \neg q\}$ $\quad\quad$ $\Delta = \{p,\ \neg p,\ q\}$

Clearly, there is a sense in which Γ is less stable, i.e. more incoherent, then Δ. More specifically, we can define a function to measure the relative *level of incoherence* of a set. By an n-covering of a set Γ, we mean a collection, $\mathcal{C} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$, of non-empty subsets of Γ such that $\Gamma = \bigcup \mathcal{C}$ (where $n \leq \omega$). Elements of an n-covering are called *clusters*. An n-covering is Σ-consistent iff each of its clusters is Σ-consistent.

Definition 2.4.1
The ℓ_Σ-value of a set Γ is defined as:

$$\ell_\Sigma(\Gamma) = \begin{cases} 0 & \textit{if } \Gamma = \emptyset \textit{ or } \Gamma \subseteq \{A : \vdash A\} \\ \\ \textit{the cardinality of the least} & \textit{if such a covering exists} \\ \Sigma\textit{-consistent covering of } \Gamma & \\ \textit{up to and including } \omega & \\ \\ \infty & \textit{otherwise} \end{cases}$$

We use $\mathfrak{C}_{\ell_\Sigma}(\Gamma)$ to denote the set of all $\ell_\Sigma(\Gamma)$-fold coverings of Γ. The sentence '$\ell_\Sigma(\Gamma) = \infty$' does not say that Γ has infinite Σ-level; rather it says that Γ has no Σ-level at all. So we must distinguish between $\ell_\Sigma(\Gamma) = \infty$ and $\ell_\Sigma(\Gamma) = \omega$. More specifically if Γ contains a Σ-villain, then $\ell_\Sigma(\Gamma) = \infty$. Also observe that if $\ell_\Sigma(\Gamma) = n \neq \infty$, then there must be a Σ-consistent n-covering of Γ.

Though the ℓ_Σ function offers us a natural way to classify inconsistent sets, it is sensitive to the syntax of the premises. Consider for instance,

Example 2.4.2
$\Sigma = \{q\}$ $\qquad \Gamma = \{p \wedge \neg p\}$ $\qquad \Delta = \{p, \neg p\}$

According to our definition, the Σ-level of Γ and the Σ-level of Δ are distinct – $\ell_\Sigma(\Gamma) = \infty$ but $\ell_\Sigma(\Delta) = 2$. However, other less syntax-sensitive means to classify inconsistent sets are available. In [80], Grant proposes three model theoretic means to classify inconsistent first order theories. To our knowledge, Grant is the first to offer such systematic classifications of inconsistent theories.

2.4.2 Quantity of Empirical Information

Some inconsistent data are less informative than others. While we agree that it is difficult to come up with a useful definition of value of information, we do not agree with Aisbett and Gibbon in [6] that inconsistent data provides no information to a decision maker. What is and what isn't informative seems to depend, at least partly, on the goal of the agent in possession of the data. For a tax auditor, inconsistencies in a taxpayer's records are useful information for detecting possible fraud. Inconsistencies may also be useful in cases where they are deployed as directives to guide learning or as indicators for faulty components in a complex system. Hence we need to develop a theoretical framework to distinguish different sorts of inconsistent data. In [123], a definition for measuring the amount of *semantic information* of an inconsistent set is given. In this section we give a definition for measuring the amount of *empirical information* in an inconsistent set.

By a quasi-model of Γ, we mean any two-valued model of any $\mathcal{A} \in M_\Sigma(\Gamma)$. Taking Γ to be a set of empirical data, i.e. data about the state of the world, we may intuitively interpret each quasi-model as representing a possible state of the world according to Γ. To define the relative quantity of empirical information of an inconsistent set, we first define the following function:

Definition 2.4.2
The λ_Σ-value of a set Γ is defined as:

$$\lambda_\Sigma(\Gamma) = \begin{cases} 0 & \textit{if } \Gamma = \emptyset \textit{ or } \Gamma \subseteq \{A : \ \vdash A\} \\ |M_\Sigma(\Gamma)| & \textit{if } M_\Sigma(\Gamma) \neq \emptyset \\ \infty & \textit{otherwise} \end{cases}$$

In effect, the λ_Σ-value is just the number of maximal Σ-consistent subsets of Γ. In terms of the relation between ℓ_Σ and λ_Σ, it is straightforward to show the following:

Proposition 2.4.1
For any $\Gamma \subseteq \Phi$, $\ell_\Sigma(\Gamma) = n \implies \lambda_\Sigma(\Gamma) \geq n$, for $1 \leq n < \omega$

Proof:
If $\ell_\Sigma(\Gamma) = n$, then there must be a Σ-consistent covering of Γ. Each cluster of such a covering of Γ is Σ-consistent and thus can be extended to a maximally Σ-consistent subset of Γ. There are n distinct and pairwise inconsistent clusters and thus there are n distinct and pairwise inconsistent extensions. Hence $n \leq \lambda_\Sigma(\Gamma)$ as required. ∎

Since Σ-villains are Σ-inconsistent and tautologies do not contribute any information about the world, we may disregard them when we are considering the amount of empirical information of a set. We let the root of Γ, $R(\Gamma)$, be the set of propositional atoms occurring in the set $\bigcup M_\Sigma(\Gamma) - \{A \in \Gamma : \vdash A\}$, i.e., $R(\Gamma)$ is the set the propositional atoms occurring in Σ-witnesses that are not tautologies. In counting the number of quasi-models of Γ, we are only concerned with the number of *equivalence classes* of quasi-models with respect to $R(\Gamma)$. So the maximum possible number of such equivalence classes is $2^{|R(\Gamma)|}$. We use $Q_{R(\Gamma)}(\Gamma)$ to denote the collection of such equivalence classes of quasi-models. We note that $|Q_{R(\Gamma)}(\Gamma)| \leq 2^{|R(\Gamma)|}$.

Definition 2.4.3

The quantity of empirical information of Γ *is given by:*

$$I_\Sigma(\Gamma) = \begin{cases} |R(\Gamma)| - \log_2 |Q_{R(\Gamma)}(\Gamma)| & \textit{if } \lambda_\Sigma(\Gamma) = 1 \\ |R(\Gamma)| - \log_2 \lambda_\Sigma(\Gamma) & \textit{if } \lambda_\Sigma(\Gamma) > 1 \\ 0 & \textit{otherwise} \end{cases}$$

When $\lambda_\Sigma(\Gamma) = 1$, $I_\Sigma(\Gamma)$ is based on the ratio between $2^{|R(\Gamma)|}$ and $|Q_{R(\Gamma)}(\Gamma)|$. When $\lambda_\Sigma(\Gamma) > 1$, $I_\Sigma(\Gamma)$ is defined by a decreasing function of the λ_Σ-value of Γ. Intuitively, the λ_Σ-value of Γ provides one possible way to measure the amount of conflict amongst the Σ-witnesses. When $\lambda_\Sigma(\Gamma) = 1$, there is no conflict and when $\lambda_\Sigma(\Gamma) > 1$ it means that there are conflicts amongst the Σ-witnesses. Moreover, the higher the λ_Σ-value of Γ, the more Σ-inconsistent subsets reside amongst the Σ-witnesses. If $\lambda_\Sigma(\Gamma) = k > 1$, then by taking the union of each distinct pair $\mathcal{A}, \mathcal{B} \in M_\Sigma(\Gamma)$ there are at least $\frac{k(k-1)}{2}$ many ways to generate Σ-inconsistent subsets amongst the Σ-witnesses. Consider the following example:

Example 2.4.3

Let $\Sigma = \{s\}$.

$\Gamma_1 =$	$\{p \wedge q,\ \neg p \wedge r,\ \neg s\}$	$\Gamma_2 =$	$\{p \wedge q \wedge r,\ \neg p \wedge q \wedge r,\ p \wedge \neg q \wedge r,\ \neg s\}$
$R(\Gamma_1)$	$\{p, q, r\}$	$R(\Gamma_2)$	$\{p, q, r\}$
$\|R(\Gamma_1)\|$	3	$\|R(\Gamma_2)\|$	3
$\lambda_\Sigma(\Gamma_1)$	2	$\lambda_\Sigma(\Gamma_2)$	3
$I_\Sigma(\Gamma_1)$	2.00	$I_\Sigma(\Gamma_2)$	1.42

Table 2.3: A comparison of two sets.

In table (2.3), $R(\Gamma_1)$ and $R(\Gamma_2)$ are identical. Moreover, since $\neg s$ is a Σ-villain s is not in $R(\Gamma_i)$, $i = 1, 2$. The λ_Σ-value of Γ_1 is lower and so the amount of conflict in the set of Σ-witnesses in Γ_1 is also lower. Consequently, $I_\Sigma(\Gamma_1) > I_\Sigma(\Gamma_2)$.

2.5 Σ-Forced Consequence

In this section we introduce a new paraconsistent consequence, called Σ-forced consequence, based on the notion of Σ-level. Σ-forced consequence is a generalization of a paraconsistent consequence operator introduced in [164; 167].

Definition 2.5.1
Σ-forced consequence: $A \in C_{F\Sigma}(\Gamma)$ *iff for each* $\mathcal{C} \in \mathfrak{C}_{\ell_\Sigma}(\Gamma)$, $\mathcal{A} \vdash A$ *for some* $\mathcal{A} \in \mathcal{C}$

Σ-forced consequence:* $A \in C^*_{F\Sigma}(\Gamma)$ *iff for each* $\mathcal{C} \in \mathfrak{C}_{\ell_\Sigma}(\Gamma)$, $\mathcal{A} \cup \Sigma \vdash A$ *for some* $\mathcal{A} \in \mathcal{C}$

where $\mathfrak{C}_{\ell_\Sigma}(\Gamma)$ *is the set of all* $\ell_\Sigma(\Gamma)$*-fold coverings of* Γ.

In other words, A is a Σ-forced consequence of Γ iff for every $\ell_\Sigma(\Gamma)$-fold covering of Γ, there is a cluster which classically implies A. Again, the main difference between $C_{F\Sigma}$ and $C^*_{F\Sigma}$ is the role Σ plays in deriving conclusions. Similarly to the previous results, any Σ-forced consequence is a Σ-forced consequence*, i.e. for any $\Gamma \subseteq \Phi$, $C_{F\Sigma}(\Gamma) \subseteq C^*_{F\Sigma}(\Gamma)$.

We also note that $C_{F\Sigma}$ and $C^*_{F\Sigma}$ are defined relative to the Σ-level of a set. Since the Σ-level of a set is not closed under supersets in general, $C_{F\Sigma}$ and $C^*_{F\Sigma}$ are both non-monotonic with respect to Γ. However, if we define Σ-forced consequence and consequence* relative to a fixed n, for $n \in \mathbb{N}$, (i.e., replace 'every $\ell_\Sigma(\Gamma)$-fold covering' with 'every n-covering' in the definition), then the resulting notions of consequence are monotonic with respect to Γ. Nonetheless, these consequences are unprincipled when $\ell_\Sigma(\Gamma) > n$. From a nonmonotonic reasoning perspective, it would be of some theoretical interest to study a varying-Σ approach to Σ-consequence. For instance, it is easy to see that for a fixed premise set Γ, if $\Sigma' \supseteq \Sigma$, then $C_{x\Sigma}(\Gamma) \subseteq C_{x\Sigma'}(\Gamma)$, where $x \in \{E, A, F, L, U, S\}$. In effect, we need to distinguish between two kinds of nonmonotonicity – those with respect to the premise set and those with respect to the constraint set. This is particularly interesting in modelling agents who are endowed with meta-beliefs that govern and provide constraints on lower level beliefs. Intriguing as it may be, however, we will not work out the details of the varying-Σ approach here.

On the assumption that a premise set Γ does not contain any Σ-villain, the relationship between Σ-forced consequence and other Σ-consequences (of Γ) is summarized in figure (2.2). Downward arrows indicate set inclusions.

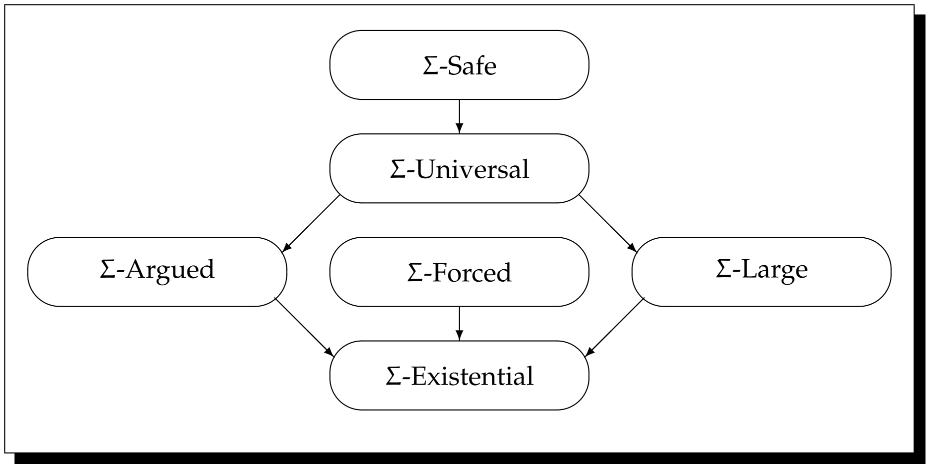

Figure 2.2: Inclusion ordering of Σ-consequences.

2.6 Preservation

In this section we will focus on the preservational properties of our inference mechanisms in terms of the ℓ_Σ, λ_Σ and I_Σ values of premise sets. We can characterise the preservational property of a consequence operator C both locally and globally. The local characterisation specifies the effect of extending the premise set by a single consequence; whereas the global characterization specifies the effect of extending the premise set by the entire consequence set. These notions are given formally in the following definitions:

Definition 2.6.1
Let C *be a consequence operator defined over the language* Φ*, i.e.,* $C : \wp(\Phi) \longrightarrow \wp(\Phi)$*. We say that* C *is*

locally ℓ_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$ *and* $A \in C(\Gamma)$, $\ell_\Sigma(\Gamma) = n$ *only if* $\ell_\Sigma(\Gamma, A) = n$

globally ℓ_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$, $\ell_\Sigma(\Gamma) = n$ *only if* $\ell_\Sigma(\Gamma, C(\Gamma)) = n$

locally λ_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$ *and* $A \in C(\Gamma)$, $\lambda_\Sigma(\Gamma) = n$ *only if* $\lambda_\Sigma(\Gamma, A) = n$

globally λ_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$, $\lambda_\Sigma(\Gamma) = n$ *only if* $\lambda_\Sigma(\Gamma, C(\Gamma)) = n$

locally I_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$ *and* $A \in C(\Gamma)$ *with* $R(A) \subseteq R(\Gamma)$, $I_\Sigma(\Gamma) = n$ *only if* $I_\Sigma(\Gamma, A) = n$

globally I_Σ-preserving *iff for any* $\Gamma \subseteq \Phi$,

1. $R(C(\Gamma)) \subseteq R(\Gamma)$

2. $\lambda_\Sigma(\Gamma) = 1$ *only if* $\lambda_\Sigma(\Gamma, C(\Gamma)) = 1$

3. $I_\Sigma(\Gamma) = n$ *only if* $I_\Sigma(\Gamma, C(\Gamma)) = n$

We note that to show that a consequence operator is not globally x-preserving, it suffices to show that it is not locally x-preserving.

Proposition 2.6.1

1. *For* $x \in \{\ell_\Sigma,\ \lambda_\Sigma,\ I_\Sigma\}$*, any globally* x*-preserving consequence operator is also locally* x*-preserving.*

2. *Let* C_1 *and* C_2 *be consequence operators such that for any* $\Gamma \subseteq \Phi$*,* $C_1(\Gamma) \subseteq C_2(\Gamma)$*. If* $C_2(\Gamma)$ *is locally (or globally)* x*-preserving, then* $C_1(\Gamma)$ *is locally (or globally)* x*-preserving for* $x \in \{\ell_\Sigma,\ \lambda_\Sigma,\ I_\Sigma\}$.

Proof:

(1) We'll assume that C is globally x-preserving and consider each case in turn:

(ℓ_Σ): We note that ℓ_Σ is a monotonically increasing function over $\subseteq$-ordering of Φ. Hence given that $\Gamma \subseteq (\Gamma \cup \{A\}) \subseteq (\Gamma \cup C(\Gamma))$ holds for any $A \in C(\Gamma)$, we have

$$\ell_\Sigma(\Gamma) \leq \ell_\Sigma(\Gamma, A) \leq \ell_\Sigma(\Gamma, C(\Gamma))$$

By the initial assumption C is globally ℓ_Σ-preserving and thus $\ell_\Sigma(\Gamma) = n$ implies that $\ell_\Sigma(\Gamma, C(\Gamma)) = n$ and hence $\ell_\Sigma(\Gamma, A) = n$ as required.
(λ_Σ): Similar to ℓ_Σ. λ_Σ is also monotonically increasing.
(I_Σ): There are two cases to consider:
($\lambda_\Sigma(\Gamma) = 1$): Then by (1) and (2) of definition (2.6.1), we have $\lambda_\Sigma(\Gamma, C(\Gamma)) = 1$ and $R(\Gamma) = R(\Gamma, A) = R(\Gamma, C(\Gamma))$. It follows that $\lambda_\Sigma(\Gamma, A)) = 1$ for every $A \in C(\Gamma)$. Hence it follows that

$$\begin{aligned} |R(\Gamma)| - \log_2 |Q_{R(\Gamma)}(\Gamma)| &= |R(\Gamma, C(\Gamma))| - \log_2 |Q_{R(\Gamma, C(\Gamma))}(\Gamma, C(\Gamma))| \\ &= |R(\Gamma, A)| - \log_2 |Q_{R(\Gamma, A)}(\Gamma, A)| \end{aligned}$$

i.e. $I_\Sigma(\Gamma) = I_\Sigma(\Gamma, A)$ as required.
($\lambda_\Sigma(\Gamma) > 1$): Then by the previous result, $\lambda_\Sigma(\Gamma) = \lambda_\Sigma(\Gamma, A)$ for each $A \in C(\Gamma)$ and $R(\Gamma) = R(\Gamma, A) = R(\Gamma, C(\Gamma))$. Hence,

$$|R(\Gamma)| - \lambda_\Sigma(\Gamma) = |R(\Gamma, A)| - \lambda_\Sigma(\Gamma, A)$$

i.e. $I_\Sigma(\Gamma) = I_\Sigma(\Gamma, A)$ as required.

(2) The cases for ℓ_Σ and λ_Σ are straightforward given that these functions are both monotonically increasing and that $C_1(\Gamma) \subseteq C_2(\Gamma)$. We'll consider the case for I_Σ:
(i) C_2 is globally I_Σ-preserving. If $\lambda_\Sigma(\Gamma) = 1$, then $\lambda_\Sigma(\Gamma, C_2(\Gamma)) = 1$ by (2) of definition (2.6.1). So by the monotonic increasing property of λ_Σ, $\lambda_\Sigma(\Gamma, C_1(\Gamma)) = 1$. By (1) of definition (2.6.1) and the assumption that $C_1(\Gamma) \subseteq C_2(\Gamma)$, we have $R(\Gamma) = R(\Gamma, C_2(\Gamma)) = R(\Gamma, C_1(\Gamma))$. Hence we have

$$\begin{aligned} |R(\Gamma)| - \log_2 |Q_{R(\Gamma)}(\Gamma)| &= |R(\Gamma, C_2(\Gamma))| - \log_2 |Q_{R(\Gamma, C_2(\Gamma))}(\Gamma, C_2(\Gamma))| \\ &= |R(\Gamma, C_1(\Gamma))| - \log_2 |Q_{R(\Gamma, C_1(\Gamma))}(\Gamma, C_1(\Gamma))| \end{aligned}$$

i.e. $I_\Sigma(\Gamma) = I_\Sigma(\Gamma, C_1(\Gamma))$ as required.
If $\lambda_\Sigma(\Gamma) > 1$, then by the monotonic increasing property of λ_Σ and the fact that $C_1(\Gamma) \subseteq C_2(\Gamma)$, we have $\lambda_\Sigma(\Gamma) = n$ implies that $\lambda_\Sigma(\Gamma, C_1(\Gamma)) = n$. In either case C_1 is globally I_Σ-preserving on the assumption that C_2 is globally I_Σ-preserving.
(ii) C_2 is locally I_Σ-preserving. Consider an arbitrary $A \in C_1(\Gamma)$ with $R(A) \subseteq R(\Gamma)$. Given $C_1(\Gamma) \subseteq C_2(\Gamma)$, $A \in C_2(\Gamma)$ follows. Hence by the local I_Σ-preserving property of C_2, $I_\Sigma(\Gamma) = n$ implies $I_\Sigma(\Gamma, A) = n$. But A was arbitrary so for any $A \in C_1(\Gamma)$ with $R(A) \subseteq R(\Gamma)$, we have $I_\Sigma(\Gamma) = n$ implies $I_\Sigma(\Gamma, A) = n$. Hence C_1 is locally I_Σ-preserving. ∎

We can summarize proposition (2.6.1) with figure (2.3). Arrows indicate implications between two statements.

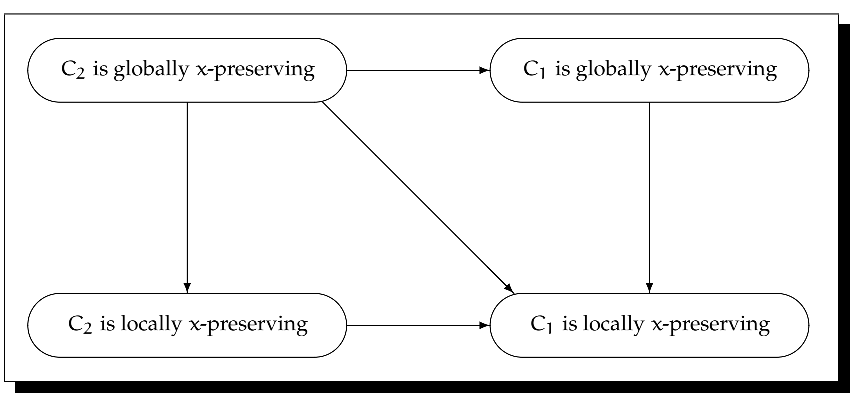

Figure 2.3: Local and global preservation for $C_1 \subseteq C_2$

Thus by propositions (2.2.4) and (2.6.1), to show that a Σ-consequence is x-preserving it suffices to show that its Σ-consequence* counterpart is x-preserving (note the exception for $C_{A\Sigma}$ and $C^{*}_{A\Sigma}$). And to show that a Σ-consequence is not globally x-preserving it suffices to show that it is not locally x-preserving. In terms of the classical consequence operator **Cn** however, it is clear that for $x \in \{\ell_{\emptyset},\ \lambda_{\emptyset},\ I_{\emptyset}\}$ **Cn** is neither locally nor globally x-preserving (since for an inconsistent Γ, $\mathbf{Cn}(\Gamma) = \Phi$).

Proposition 2.6.2

1. *For every $x \in \{E\Sigma,\ A\Sigma,\ F\Sigma,\ L\Sigma,\ U\Sigma,\ S\Sigma\}$, $C^{(*)}_{x}$ fails to be globally I_{Σ}-preserving.*
2. *For every $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma},\ I_{\Sigma}\}$, $C_{A\Sigma}$ fails to be globally and locally x-preserving.*
3. *For every $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma},\ I_{\Sigma}\}$, $C^{(*)}_{E\Sigma}$ fails to be globally and locally x-preserving.*
4. *For every $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma},\ I_{\Sigma}\}$, $C^{*}_{A\Sigma}$ is locally x-preserving but not globally x-preserving.*
5. *For $x \in \{\lambda_{\Sigma},\ I_{\Sigma}\}$, $C^{(*)}_{F\Sigma}$ is not globally and locally x-preserving, but are globally and locally ℓ_{Σ}-preserving.*
6. *For $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma},\ I_{\Sigma}\}$, $C^{(*)}_{L\Sigma}$ fail to be globally and locally x-preserving.*
7. *For $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma}\}$, $C^{(*)}_{U\Sigma}$ and $C^{(*)}_{S\Sigma}$ are locally and globally x-preserving. Moreover, $C^{(*)}_{U\Sigma}$ and $C^{(*)}_{S\Sigma}$ are locally I_{Σ}-preserving but not globally I_{Σ}-preserving.*

Proof:

(1) It suffices to observe that $A \vdash A \vee B$ is a valid classical rule and thus using right weakening if $A \in C^{(*)}_{x}(\Gamma)$, then $A \vee B \in C^{(*)}_{x}(\Gamma)$ for any B. Hence $R(C^{(*)}_{x}(\Gamma)) \not\subseteq R(\Gamma)$.

(2) It suffices to show that $C_{A\Sigma}$ is not locally x-preserving for $x \in \{\ell_{\Sigma},\ \lambda_{\Sigma},\ I_{\Sigma}\}$. We give counterexamples for each case:

(ℓ_{Σ}): Clearly $s \wedge t \in C_{A\Sigma}(\Gamma)$ for the following Σ and Γ:

$$\begin{array}{ll} \Sigma = \{\neg s \vee \neg t \vee \neg q,\ \neg s \vee \neg t \vee \neg r\} & \ell_{\Sigma}(\Gamma) = 2 \\ \Gamma = \{p \wedge q,\ \neg p \wedge r,\ s,\ t\} & \ell_{\Sigma}(\Gamma, s \wedge t) = 3 \end{array}$$

(λ_{Σ}, I_{Σ}): Clearly $p \vee (\neg s \vee \neg t) \in C_{A\Sigma}(\Gamma)$ for the following Σ and Γ:

$$\begin{array}{ll} \Sigma = \emptyset & \lambda_{\Sigma}(\Gamma) = 2 \\ \Gamma = \{p,\ \neg p \wedge s,\ \neg p \wedge t\} & I_{\Sigma}(\Gamma) = 2 \\ \lambda_{\Sigma}(\Gamma, p \vee (\neg s \vee \neg t)) = 4 & I_{\Sigma}(\Gamma, p \vee (\neg s \vee \neg t)) = 1 \end{array}$$

(3) It follows from propositions (2.2.2), (2.6.1) and (2) above.

(4) We consider each case in turn:

(ℓ_Σ): Let $A \in C^*_{A\Sigma}(\Gamma)$ and let $\mathcal{C} = \{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ be a Σ-consistent covering of Γ that witnesses $\ell_\Sigma(\Gamma) = n$. We claim that for each $i \leq n$, $\mathcal{C}_i \cup \{A\}$ is Σ-consistent and thus $\Big\{\mathcal{C}_1 \cup \{A\}, \ \ldots, \ \mathcal{C}_n \cup \{A\}\Big\}$ is a Σ-consistent covering of $\Gamma \cup \{A\}$. The claim clearly holds since for each $i \leq n$, $\mathcal{C}_i$ is contained in some $\mathcal{A} \in M_\Sigma(\Gamma)$ and for every $\mathcal{A} \in M_\Sigma(\Gamma)$, $\Sigma \cup \mathcal{A} \nvdash \neg A$. So for each $i \leq n$, $\Sigma \cup \mathcal{C}_i \nvdash \neg A$. By the minimality of n, $\ell_\Sigma(\Gamma, A) = n$ as required.

(λ_Σ): similar to the ℓ_Σ case.

(I_Σ): Consider any $A \in C^*_{A\Sigma}(\Gamma)$ with $R(A) \subseteq R(\Gamma)$. There are two cases to consider:

($\lambda_\Sigma(\Gamma) = 1$): Then by ordinary classical logic the following equality holds:

$$|Q_{R(\Gamma)}(\Gamma)| = |Q_{R(\Gamma, A)}(\Gamma, A)|$$

So $I_\Sigma(\Gamma) = I_\Sigma(\Gamma, A)$ as required.

($\lambda_\Sigma(\Gamma) = n > 1$): Then by the previous result $C^*_{A\Sigma}$ is λ_Σ-preserving and thus

$$|R(\Gamma)| - \log_2 \lambda_\Sigma(\Gamma) = |R(\Gamma, A)| - \log_2 \lambda_\Sigma(\Gamma, A)$$

Hence $I_\Sigma(\Gamma) = I_\Sigma(\Gamma, A)$ as required. Failure of global preservation for $C^*_{A\Sigma}$ follows immediately from (3) of proposition (2.2.4).

(5) For ℓ_Σ-preservation, it suffices to show that $C^*_{F\Sigma}$ is globally ℓ_Σ-preserving. Assume that $\ell_\Sigma(\Gamma) = n$. We note that by the minimality of n, $\ell_\Sigma(\Gamma, C^*_{F\Sigma}(\Gamma)) \geq n$. Consider an arbitrary but fixed Σ-consistent covering of Γ, $\mathcal{C} = \{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$. Let

$$\mathcal{C}^* = \Big\{\mathbf{Cn}(\Sigma \cup \mathcal{C}_i) : \ i \leq n\Big\}$$

We claim that

(a) $\Gamma \cup C^*_{F\Sigma}(\Gamma) \subseteq \bigcup \mathcal{C}^*$: If $A \in \Gamma$, then clearly $A \in \bigcup \mathcal{C}^*$ since $\mathcal{C}$ is an n-covering of Γ. If $A \in C^*_{F\Sigma}(\Gamma)$, then every n-covering of Γ contains a cluster which together with Σ classically implies A. In particular for $\mathcal{C}$, there must be a $j \leq n$ such that $\Sigma \cup \mathcal{C}_j \vdash A$, i.e. $A \in \mathbf{Cn}(\Sigma \cup \mathcal{C}_j)$. Hence $A \in \bigcup \mathcal{C}^*$.

(b) $\ell_\Sigma(\bigcup \mathcal{C}^*) = n$: trivial given that for each $i \leq n$, $\mathcal{C}_i$ is Σ-consistent.

(c) $\ell_\Sigma(\Gamma, C^*_{F\Sigma}(\Gamma)) = n$: from (a) and (b) above, we have $\Gamma \subseteq \Big(\Gamma \cup C^*_{F\Sigma}(\Gamma)\Big) \subseteq \bigcup \mathcal{C}^*$. But $\ell_\Sigma(\Gamma) = \ell_\Sigma(\bigcup \mathcal{C}^*) = n$, hence $\ell_\Sigma(\Gamma, C^*_{F\Sigma}(\Gamma)) = n$.

To verify that $C^{(*)}_{F\Sigma}$ is neither λ-preserving nor I_Σ-preserving, it suffices to consider the following example where $(p \vee \neg q) \vee (\neg r \vee \neg s) \in C_{F\Sigma}(\Gamma)$:

$$\Sigma = \emptyset \qquad \lambda_\Sigma\Big(\Gamma, (p \vee \neg q) \vee (\neg r \vee \neg s)\Big) = 6$$

$$\Gamma = \{p,\ q,\ \neg q, \neg p \wedge r,\ \neg p \wedge s\} \qquad \lambda_\Sigma(\Gamma) = 4$$

(6) It suffices to show that for $x \in \{\ell_\Sigma,\ \lambda_\Sigma, I_\Sigma\}$ $C_{L\Sigma}$ fails to be locally x-preserving. We consider each case in turn:

(ℓ_Σ): We can modify the example used in (2). Clearly $s \wedge t \in C_{L\Sigma}(\Gamma)$ for the following Σ and Γ:

$$\Sigma = \{\neg s \vee \neg t \vee \neg q,\ \neg s \vee \neg t \vee \neg r\} \qquad \ell_\Sigma(\Gamma, s \wedge t) = 3$$

$$\Gamma = \{u \wedge p \wedge q,\ u \wedge \neg p \wedge r,\ \neg u \wedge s,\ \neg u \wedge t\} \qquad \ell_\Sigma(\Gamma) = 2$$

(λ_Σ, I_Σ): Again a modification of the example used in (2) suffices. Clearly $p \vee (\neg s \vee \neg t) \in C_{L\Sigma}(\Gamma)$ for the following Σ and Γ:

$$\Sigma = \emptyset \qquad I_\Sigma(\Gamma, p \vee (\neg s \vee \neg t)) = 3$$

$$\Gamma = \{p,\ q,\ r,\ \neg p \wedge \neg q \wedge \neg r \wedge s,\ \neg p \wedge \neg q \wedge \neg r \wedge t\} \qquad \lambda_\Sigma(\Gamma) = 2$$

$$I_\Sigma(\Gamma) = 4 \qquad \lambda_\Sigma(\Gamma, p \vee (\neg s \vee \neg t)) = 4$$

(7) The first part of the statement is straightforward since we have

$$C^*_{U\Sigma}(\Gamma) = \bigcap_{\mathcal{A} \in M_\Sigma(\Gamma)} \mathbf{Cn}\Big(\Sigma \cup \mathcal{A}\Big)$$

The second part of the statement is also straightforward since for any $A \in C^*_{U\Sigma}(\Gamma)$ with the property that $R(A) \subseteq R(\Gamma)$, $I_\Sigma(\Gamma)$ and $I_\Sigma(\Gamma, A)$ are clearly equal. ∎

The preservational properties of our Σ-consequences and Σ-consequences* are summarized in table (2.4). '+' ('−') indicates that the relevant property is (is not) preserved.

	local			global		
	ℓ_Σ	λ_Σ	I_Σ	ℓ_Σ	λ_Σ	I_Σ
$C_{E\Sigma}$	–	–	–	–	–	–
$C_{A\Sigma}$	–	–	–	–	–	–
$C_{F\Sigma}$	+	–	–	+	–	–
$C_{L\Sigma}$	–	–	–	–	–	–
$C_{U\Sigma}$	+	+	+	+	+	–
$C_{S\Sigma}$	+	+	+	+	+	–

	local			global		
	ℓ_Σ	λ_Σ	I_Σ	ℓ_Σ	λ_Σ	I_Σ
$C^*_{E\Sigma}$	–	–	–	–	–	–
$C^*_{A\Sigma}$	+	+	+	–	–	–
$C^*_{F\Sigma}$	+	–	–	+	–	–
$C^*_{L\Sigma}$	–	–	–	–	–	–
$C^*_{U\Sigma}$	+	+	+	+	+	–
$C^*_{S\Sigma}$	+	+	+	+	+	–

Table 2.4: Preservational properties of Σ-consequences.

2.6.1 Maximality

Since for each consequence operator C we can define a consequence relation $\vdash_C$ such that $\langle \Gamma, A\rangle \in \vdash_C$ iff $A \in C(\Gamma)$, we may speak of the consequence relation $\vdash_C$ as being induced by C. Furthermore we say that $\vdash_C$ is (locally or globally) x-preserving iff C is. One important fact is that strictly speaking there is no smallest (locally or globally) x-preserving consequence relation, $x \in \{\ell_\Sigma, \lambda_\Sigma, I_\Sigma\}$. By this we mean that for a fixed x the intersection of all x-preserving consequence relations (induced by their respective consequence operators) is in fact empty. However, it is possible that two consequence operators C_1 and C_2 may be related in such a way that (1) C_1 is (locally or globally) x-preserving but C_2 is not, and (2) for any Γ, $C_1(\Gamma)$ is contained in $C_2(\Gamma)$. In such a case it is natural to ask whether $\vdash_{C_1}$ can be extended maximally within $\vdash_{C_2}$ to a (locally or globally) x-preserving consequence relation. In fact, this is exactly the situation at hand. For instance, $C_{U\Sigma}$ is globally ℓ_Σ-preserving but $C_{L\Sigma}$ is not (moreover for any Γ, $C_{U\Sigma}(\Gamma) \subseteq C_{L\Sigma}(\Gamma)$). So a natural question is whether $\vdash_{C_{U\Sigma}}$ can be extended maximally to a ℓ_Σ-preserving extension within $\vdash_{C_{L\Sigma}}$. Such maximal extensions are theoretically interesting since they allow us to deduce more conclusions while still preserving the relevant property in question.

2.6.2 Special Conditions

Another theoretically interesting question is whether there are special conditions under which a particular inference mechanism can preserve a property even though the mechanism does not preserve the property in general. We may think of these special conditions as *application* conditions which allow us to use certain inference mechanisms to preserve certain properties. For instance, if each maximal Σ-consistent subset of a premise set Γ has the same cardinality, then $C_{L\Sigma}(\Gamma)$ is identical to $C_{U\Sigma}(\Gamma)$. So the λ_Σ value of Γ is preserved by $C_{L\Sigma}$ in this case even though $C_{L\Sigma}$ is neither locally nor globally λ_Σ-preserving in general. For instance, the following fact allows us to use $C^*_{A\Sigma}$ to globally preserve the ℓ_Σ value of Γ when $\lambda_\Sigma(\Gamma) = n < \omega$.

Proposition 2.6.3
For any $\Gamma \subseteq \Phi$, $\ell_\Sigma(\Gamma) = \lambda_\Sigma(\Gamma) = n < \omega \implies \ell_\Sigma(\Gamma, C^*_{A\Sigma}(\Gamma)) = n$.

Proof:
We'll assume that Γ is arbitrary and that $\ell_\Sigma(\Gamma) = \lambda_\Sigma(\Gamma) = n < \omega$. We construct a Σ-consistent n-covering of $\Gamma \cup C^*_{A\Sigma}(\Gamma)$. Let $M_\Sigma(\Gamma) = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$. Let $\mathcal{C} = \{\mathcal{B}_1, \ldots, \mathcal{B}_n\}$ be an arbitrary but fixed Σ-consistent n-covering of Γ where the enumeration of $\mathcal{C}$ is such that for each $i \leq n$, $\mathcal{B}_i \subseteq \mathcal{A}_i$. This is clearly possible since every cluster in $\mathcal{C}$ is Σ-consistent. For each $\mathcal{A}_i \in M_\Sigma(\Gamma)$ we define:

$$\mathcal{A}^*_i = \left\{ A \in \mathbf{Cn}(\Sigma \cup \mathcal{A}_i) : \neg A \notin \bigcup_{j=1}^{n} \mathbf{Cn}(\Sigma \cup \mathcal{A}_j) \right\}$$

Clearly, $\mathcal{A}^*_i \subseteq \mathbf{Cn}(\Sigma \cup \mathcal{A}_i)$. Hence each $\mathcal{A}^*_i$ is Σ-consistent. It is straightforward to verify that

$$\bigcup_{i=1}^{n} \mathcal{A}^*_i = C^*_{A\Sigma}(\Gamma)$$

We now define

$$\mathcal{C}' = \{\mathcal{B}_1 \cup \mathcal{A}^*_1, \ldots, \ \mathcal{B}_n \cup \mathcal{A}^*_n\}$$

$\mathcal{C}'$ is clearly a Σ-consistent n-covering of $\Gamma \cup C^*_{A\Sigma}(\Gamma)$. Since $\Gamma \subseteq \Gamma \cup C_{A\Sigma}(\Gamma)$ and $\ell_\Sigma(\Gamma) = n$, by the minimality of n and the monotonic increasing property of ℓ_Σ, $\ell_\Sigma(\Gamma, C^*_{A\Sigma}(\Gamma)) = n$ as required. ∎

2.6.3 Combining Inference Mechanisms

Finally, we have not considered the effect of combining different inference mechanisms. For instance by taking the union and intersection of $\vdash_{C_{F\Sigma}}$ and $\vdash_{C_{U\Sigma}}$ we can obtain two new consequence relations. Clearly, $\vdash_{C_{F\Sigma}} \cap \vdash_{C_{U\Sigma}}$ is both non-empty and ℓ_Σ-preserving (since $C_{F\Sigma}$ and $C_{U\Sigma}$ are both ℓ_Σ-preserving). Again from the point of view of section (2.6.1), $\vdash_{C_{F\Sigma}} \cup \vdash_{C_{U\Sigma}}$ is a more interesting option since it extends both $\vdash_{C_{F\Sigma}}$ and $\vdash_{C_{U\Sigma}}$.

2.7 Conclusion

In this chapter we have applied a preservation-theoretic approach to analyze and compare six different inconsistency tolerant inference mechanisms. The crux of our motivation is to demonstrate that *truth is not the only property worthy of preservation*. Which properties are to be preserved in an inference can depend on our interests and goals. As the late Jon Barwise puts it:

> ...the study of valid inference as a situated activity shifts attention from *truth preservation to information extraction* and *information processing*. Valid inference is seen not as a relation between sentences that simply preserves truth, but rather as a situated, purposeful activity whose aim is the extraction of information from a situation, information relevant to the agent. ([16], p.xiv)

In a broader context, the notion of preservation can provide a theoretically rich framework for understanding a variety of formalisms. In future work, we hope to extend our approach to analyse belief revision mechanisms, nonmonotonic reasoning systems and other practical reasoning systems.

Chapter 3

Rescher-Mechanism

3.1 Introduction

A common complaint against reasoning based on maximal consistent subsets is that it is too sensitive to the underlying syntax of the logical representation. This may result in information being isolated, and thereby preventing useful information to be extracted. Consider the following example:

Example 3.1.1
Two information sources may disagree with respect to p while not disagreeing in other respects: $\Gamma = \{p \wedge \neg q, \neg p \wedge (q \vee r)\}$

In our example, there is a sufficiently clear sense in which neither $\neg q$ nor $q \vee r$ are directly involved in an inconsistency, though they are conjoined with something that is inconsistent. Splitting Γ into consistent subsets will prevent us from deducing the potentially useful information r. Hence according to Belnap,

> ...Rescher's method gives wildly different accounts depending on just how many ampersands are replaced by commas, or vice versa. It depends too much on how our ...subtheory ...is itself separated into sentential bits. (page 544 [8])

Belnap's criticism is fair. It is intuitively implausible that an inference mechanism for handling inconsistent information should give wildly different conclusions for minor syntactic variations in the logical representation. But Belnap's criticism also applies to other formal mechanisms for handling inconsistency such as belief revision. A *syntax-based* revision of Γ with r would require us to give up at least one member of Γ (see Nebel's [131; 132]). This may incur unwanted information loss.

In 1979 [23], Belnap proposed a particular amendment to Rescher's strategy for reasoning with maximal consistent subsets. In 1989 [24] Belnap changed his mind and made a further amendment to his earlier amendment. More recently in [87], Horty explicitly endorsed Belnap's second amendment to address a related problem in handling inconsistent instructions and commands. In actual fact, Belnap's suggestions on both occasions amount to the same strategy of finding different ways to *articulate* the input logical representation. According to Belnap, the input logical description Γ is first to be closed under some non-classical logic generating a superset Γ^* and then Rescher's strategy can be applied to Γ^* in the normal way. The role of Γ^* is to make explicit the content of Γ so that the kind of difficulties that arise in situations similar to example (3.1.1) can be avoided. In 1979, Belnap's suggestion was to use Angell's analytic containment. In 1989 [24], Belnap's suggestion was to use an even more restrictive non-classical logic based on the idea of *conjunctive containment*.

In this chapter, we'll first highlight the connection between Rescher's method of reasoning from maximal consistent subsets and the default reasoning of Reiter [151]. This gives us the necessary background to appreciate Belnap's criticism in relation to more recent developments in AI. This also gives us a direct way to apply the preservation analysis from the previous chapter to various forms of default reasoning. Finally, we'll address Belnap's criticism by pointing out that his suggestion of using conjunctive containment seems to be open to the very objection he raised. We'll suggest a strategy to amend conjunctive containment.

3.2 Connection With Default Reasoning

In many ways Rescher's method of reasoning from maximal consistent subsets has anticipated many recent developments in AI. One in particular is the nonmonotonic formalism developed by Reiter in [151]. In this section, we'll recap the connection between Rescher's method and Reiter's original formalism for default reasoning. Since the publication of [151], Reiter's formalism has been revised and extended (see Schaub [163] for a summary). Many of these new developments have been shown to be expressively equivalent to various forms of belief revision formalism. We'll not be able to summarise all these new developments here. But since many of these extensions are theoretically grounded in some form of reasoning from maximal consistent subsets, Belnap's methodological criticism is still in force here. In any case, we'll focus on the standard default formalism of Reiter. We'll begin by recalling some standard definitions.

In Reiter's formalism, a default theory is a pair $\langle \mathcal{D}, \mathcal{F} \rangle$, where $\mathcal{F}$ is a set of formulae called facts and $\mathcal{D}$ is a set of default rules of the form:

$$\frac{A : B_1, \ldots, B_n}{C}$$

Intuitively, the meaning of the rule is that if A is provable and each B_i is consistent, then we may conclude C. A is called a *prerequisite* of the default rule, each B_i is a *justification* of the default rule and C is a *consequent* of the default rule. A default rule without justification is equivalent to an inference rule in standard logics. Hence the requirement of justification is responsible for the nonmonotonicity of default reasoning.

Example 3.2.1

$$\mathcal{F} = \{\texttt{head light on}\} \qquad \mathcal{F}' = \{\texttt{head light on}, \texttt{brake light fails}\}$$

$$\mathcal{D} = \left\{ \frac{\texttt{head light on: component 1 ok, ..., component n ok}}{\texttt{electrical system ok}} \right\}$$

From the observed fact that the head light is working properly and *in the absence of evidence to the contrary* it is reasonable to conclude that the electrical system is working properly. But of course this conclusion can be defeated when we observe the failure of some other component of the system even though the head light may still be in operation. So in this sense a default rule is defeasible when new observation violates some of the justifications of the rule.

The key concept in default reasoning is the notion of an extension. An extension of a default theory is a deductive closure (under classical logic) of the facts together with the consequents of the *applicable* default rules. Intuitively we can think of an extension as a possible scenario according to the information provided by the facts and the applicable default rules. More formally an extension for a default theory is defined as follows:

Definition 3.2.1
(Reiter [151]) Let $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle$ be a default theory. For any deductively closed set $\mathcal{S} \subseteq \Phi$, let $\gamma(\mathcal{S})$ be the least set satisfying the following conditions:

1. $\mathcal{F} \subseteq \gamma(\mathcal{S})$

2. $\mathbf{Cn}(\gamma(\mathcal{S})) = \gamma(\mathcal{S})$ *(where* **Cn** *is closure under classical deduction)*

3. *If* $\frac{A:B_1,\ldots,B_n}{C} \in \mathcal{D}$*,* $A \in \gamma(\mathcal{S})$*, and* $\neg B_1, \ldots, \neg B_n \notin \mathcal{S}$*, then* $C \in \gamma(\mathcal{S})$

A deductively closed set $\mathcal{E} \subseteq \Phi$ is an extension for $\mathcal{W}$ ($\mathcal{E} \in \mathtt{ext}(\mathcal{W})$) iff $\gamma(\mathcal{E}) = \mathcal{E}$, i.e. $\mathcal{E}$ is a fixed point of the operator γ.

Reiter's definition of an extension is not recursive since it appeals to a fixed point construction of γ. In example (3.2.1), the extensions of $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle$ and $\mathcal{W}' = \langle \mathcal{D}, \mathcal{F}' \rangle$ are respectively:

$$\mathtt{ext}(\mathcal{W}) = \Big\{ \mathbf{Cn}(\{\mathtt{head\ light\ on},\ \mathtt{electrical\ system\ ok}\}) \Big\}$$

$$\mathtt{ext}(\mathcal{W}') = \Big\{ \mathbf{Cn}(\{\mathtt{head\ light\ on},\ \mathtt{brake\ light\ fails}\}) \Big\}$$

However, the existence of an extension is not guaranteed in general for every default theory. Multiple extensions for a given default theory are also possible. There are three basic decision problems associated with default reasoning:

Extension Existence: Given a default theory $\mathcal{W}$, is $\mathtt{ext}(\mathcal{W})$ non-empty?

Intersection Membership: Given a default theory $\mathcal{W}$ and a formula A, is A a member of $\bigcap \mathtt{ext}(\mathcal{W})$?

Union Membership: Given a default theory $\mathcal{W}$ and a formula A, is A a member of $\bigcup \mathtt{ext}(\mathcal{W})$?

The first order version of default theory is clearly not semi-decidable with respect to these decision problems. Even in the propositional case, the complexities of these problems are generally very high – Σ_2^P and Π_2^P hard. Thus it is often desirable to identify subclasses of default theories with either guaranteed existence of extension(s) or with low computational complexity. A normal default theory for instance is one whose default rules are of the form:

$$\frac{A : B}{B}$$

For the class $\mathcal{N}$ of *normal* propositional default theories the existence of an extension is guaranteed but the complexity of determining whether a formula is in an extension is Σ_2^P complete. Analogously, the decision problem of determining whether a formula is in every extension is Π_2^P complete (see chapter 4 of [43] for a summary of complexity results for default reasoning).

Another important class of default theories is the class of *prerequisite free* default theories. A prerequisite free default theory is one whose default rules are of the form:

$$\frac{: B_1, \ldots, B_n}{C}$$

Two default theories $\mathcal{W}$ and $\mathcal{W}'$ are said to be extension equivalent if they have the same set of extensions, i.e. $\mathrm{ext}(\mathcal{W}) = \mathrm{ext}(\mathcal{W}')$. The class of prerequisite free default theories is *representationally complete* with respect to the class of all default theories in the sense that every default theory is extension equivalent to some prerequisite free default theory. A default theory $\mathcal{W}$ is said to be inconsistent if its only extension is Φ. In [151] Reiter shows that:

Theorem 3.2.1

(Reiter [151])

1. *Let Γ be a set of formulae and $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle$ be a default theory. Define $\Gamma_0 = \mathcal{F}$ and for each $i \geq 0$,*

$$\Gamma_{i+1} = \mathbf{Cn}(\Gamma_i) \cup \left\{ C : \ \frac{A : B_1, \ldots, B_n}{C} \in \mathcal{D}, \ A \in \Gamma_i, \ \neg B_1, \ldots, \neg B_n \notin \Gamma \right\}$$

Then Γ is an extension for $\mathcal{W}$ iff

$$\Gamma = \bigcup_{i=0}^{\infty} \Gamma_i$$

2. *If $\mathcal{E}$ and $\mathcal{E}'$ are extensions of a default theory and $\mathcal{E}' \subseteq \mathcal{E}$, then $\mathcal{E}' = \mathcal{E}$*

3. *A default theory $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle$ is inconsistent iff $\mathcal{F}$ is classically inconsistent iff $\mathrm{ext}(\mathcal{W}) = \{\Phi\}$.*

4. *Extensions exist for every normal default theory.*

5. *Extensions of a normal default theory are orthogonal in the sense that they are pairwise inconsistent.*

We note that in (1) of theorem (3.2.1), the definition of Γ_{i+1} makes essential reference to Γ and hence it is not a recursive definition. Using Reiter's results, we can show that there is a close relationship between *consistent prerequisite free normal* default theories and reasoning from maximal consistent subsets. Note that by (3) of theorem (3.2.1), inconsistent default theories are extension equivalent trivially. By (4) of theorem (3.2.1), extension is guaranteed to exist for any default theory in the class $\mathcal{P}$ of consistent prerequisite free normal default theories.

Theorem 3.2.2
For every $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle \in \mathcal{P}$*,* $\mathrm{ext}(\mathcal{W})$ *is of the form*

$$\{\mathbf{Cn}(\Sigma \cup \mathcal{A}) : \ \mathcal{A} \in M_{\Sigma}(\Gamma)\}$$

for some Γ *and consistent* Σ*.*

Proof:
Let $\mathcal{W} = \langle \mathcal{D}, \mathcal{F} \rangle \in \mathcal{P}$ be arbitrary but fixed. We set $\Sigma = \mathcal{F}$. First we observe that by part (3) of theorem (3.2.1), Σ is consistent and each extension of $\mathcal{W}$ is also consistent. Define

$$\Gamma = \left\{ A : \ \frac{: A}{A} \in \mathcal{D} \right\}$$

Let $\mathcal{E} \in \mathrm{ext}(\mathcal{W})$ be an arbitrary but fixed extension, consider $\Delta = \mathbf{Cn}((\mathcal{E} \cap \Gamma) \cup \Sigma)$. We claim that

Claim: $\Delta = \mathcal{E}$:

Proof of Claim: ($\Delta \subseteq \mathcal{E}$): Since $\mathcal{E} \in \mathrm{ext}(\mathcal{W})$, $\gamma(\mathcal{E}) = \mathcal{E}$ and thus by properties (1) and (2) of definition (3.2.1) we have:

$$\begin{aligned} \mathcal{E} \cap \Gamma \subseteq \mathcal{E} &\implies (\mathcal{E} \cap \Gamma) \cup \Sigma \subseteq \mathcal{E} \\ &\implies \mathbf{Cn}((\mathcal{E} \cap \Gamma) \cup \Sigma) \subseteq \mathbf{Cn}(\mathcal{E}) \\ &\implies \mathbf{Cn}((\mathcal{E} \cap \Gamma) \cup \Sigma) \subseteq \mathcal{E} \end{aligned}$$

($\Delta \supseteq \mathcal{E}$): Now we use part (1) of theorem (3.2.1). Since $\mathcal{E} \in \mathrm{ext}(\mathcal{W})$, we have $\mathcal{E} = \bigcup_{i=0}^{\infty} \mathcal{E}_i$. We prove inductively that $\mathcal{E}_i \subseteq \Delta$. For the basis it is trivial since $\mathcal{E}_0 = \Sigma \subseteq \Delta$. Now we assume the induction hypothesis that $\mathcal{E}_i \subseteq \Delta$ and prove that $\mathcal{E}_{i+1} \subseteq \Delta$. If $A \in \mathcal{E}_{i+1}$, then $A \in \mathcal{E}_i$ or $A \in \{C : \frac{:C}{C} \in \mathcal{D}, \neg C \notin \mathcal{E}\}$. In the former case, we have $A \in \Delta$ by the induction hypothesis. In the later case, we have $A \in \Gamma$ and $A \in \gamma(\mathcal{E})$ by (3) of definition (3.2.1) hence $A \in \mathcal{E}$ as $\gamma(\mathcal{E}) = \mathcal{E}$. Hence $A \in \mathcal{E} \cap \Gamma$ and so $A \in \Delta$ as required. This completes the proof that $\mathcal{E} = \bigcup_{i=0}^{\infty} \mathcal{E}_i \subseteq \Delta$.

To continue with the main proof we need to show that $\mathcal{E} \cap \Gamma \in M_{\Sigma}(\Gamma)$. Clearly, $\mathcal{E} \cap \Gamma$ is Σ-consistent since $\mathcal{E}$ is Σ-consistent. To show maximal consistency, consider any $A \in \Gamma$ such that $A \notin \mathcal{E} \cap \Gamma$. Since $\gamma(\mathcal{E}) = \mathcal{E}$ by part (3) of definition (3.2.1), we have $A \in \mathcal{E}$ or $\neg A \in \mathcal{E}$. Since $A \notin \mathcal{E} \cap \Gamma$, we have $A \notin \mathcal{E}$. Thus, $\neg A \in \mathcal{E}$. But $\Delta = \mathcal{E}$ so $(\mathcal{E} \cap \Gamma) \cup \Sigma \vdash \neg A$. Hence $(\mathcal{E} \cap \Gamma) \cup \{A\}$ is not Σ-consistent. This suffices to show that

$\mathcal{E} \cap \Gamma \in M_\Sigma(\Gamma)$. Since $\mathcal{E}$ is arbitrary, we conclude that every extension of $\mathcal{W}$ must be of the form $\mathbf{Cn}(\Sigma \cup \mathcal{A})$ where $\mathcal{A} \in M_\Sigma(\Gamma)$. Hence, $ext(\mathcal{W}) \subseteq \{\mathbf{Cn}(\Sigma \cup \mathcal{A}) : \mathcal{A} \in M_\Sigma(\Gamma)\}$.

Now to show that for every $\mathcal{A} \in M_\Sigma(\Gamma)$, $\mathbf{Cn}(\Sigma \cup \mathcal{A}) \in ext(\mathcal{W})$, we make use of part (1) of theorem (3.2.1). We consider an arbitrary $\mathcal{A} \in M_\Sigma(\Gamma)$ and define $\mathcal{S}_0 = \Sigma = \mathcal{F}$ and for each $i \geq 0$

$$\mathcal{S}_{i+1} = \mathbf{Cn}(\mathcal{S}_i) \cup \left\{ A : \frac{:A}{A} \in \mathcal{D}, \neg A \notin \mathbf{Cn}(\Sigma \cup \mathcal{A}) \right\}$$

We note that for $i \geq 2$, $\mathcal{S}_i = \mathcal{S}_{i+1}$. Hence we only need to verify that $\mathcal{S}_2 = \mathbf{Cn}(\Sigma \cup \mathcal{A})$, i.e. we need to verify that

$$\mathbf{Cn}(\Sigma \cup \mathcal{A}) \quad = \quad \mathbf{Cn}\left(\mathbf{Cn}(\Sigma) \cup \left\{ A : \frac{:A}{A} \in \mathcal{D}, \neg A \notin \mathbf{Cn}(\Sigma \cup \mathcal{A}) \right\} \right) \tag{3.1}$$

We claim that in equation (3.1),

$$\left\{ A : \frac{:A}{A} \in \mathcal{D}, \neg A \notin \mathbf{Cn}(\Sigma \cup \mathcal{A}) \right\} = \mathcal{A}$$

To verify our claim consider an arbitrary $B \in \mathcal{A}$. Since $\mathcal{A} \subseteq \Gamma$, we have $B \in \Gamma$ and thus $\frac{:B}{B} \in \mathcal{D}$. By the maximal Σ-consistency of $\mathcal{A}$, we have $\neg B \notin \mathbf{Cn}(\Sigma \cup \mathcal{A})$. Conversely consider an arbitrary $B \in \{A : \frac{:A}{A} \in \mathcal{D}, \neg A \notin \mathbf{Cn}(\Sigma \cup \mathcal{A})\}$. It follows that $B \in \Gamma$ and $\neg B \notin \mathbf{Cn}(\Sigma \cup \mathcal{A})$. Suppose to the contrary that $B \notin \mathcal{A}$, then by the maximal Σ-consistency of $\mathcal{A}$ we have $\neg B \in \mathbf{Cn}(\Sigma \cup \mathcal{A})$ contradicting our previous claim. Hence $B \in \mathcal{A}$ as required.

Thus equation (3.1) can be rewritten as:

$$\mathbf{Cn}(\Sigma \cup \mathcal{A}) \quad = \quad \mathbf{Cn}\big(\mathbf{Cn}(\Sigma) \cup \mathcal{A}\big) \tag{3.2}$$

To verify (3.2), we note that by reflexivity $\Sigma \cup \mathcal{A} \subseteq \mathbf{Cn}(\Sigma) \cup \mathcal{A}$ and thus by the monotonicity we get $\mathbf{Cn}(\Sigma \cup \mathcal{A}) \subseteq \mathbf{Cn}(\mathbf{Cn}(\Sigma \cup \mathcal{A}))$. Conversely by reflexivity and monotonicity, $\mathbf{Cn}(\Sigma) \cup \mathcal{A} \subseteq \mathbf{Cn}(\Sigma \cup \mathcal{A})$ and thus $\mathbf{Cn}\big(\mathbf{Cn}(\Sigma) \cup \mathcal{A})\big) \subseteq \mathbf{Cn}\big(\mathbf{Cn}(\Sigma \cup \mathcal{A}\big) = \mathbf{Cn}(\Sigma \cup \mathcal{A})$ by idempotence of $\mathbf{Cn}$. ∎

We note that for any given set Γ and any given consistent set Σ, the family of closures of Σ-maximal consistent subsets of Γ is precisely the family of extensions of some consistent prerequisite free normal default theory. The corresponding default theory is defined in the obvious way. It is easy to see that the following is an immediate corollary of theorem (3.2.2).

Corollary 3.2.1
Let $\mathcal{W}$ and Γ be defined as in theorem (3.2.2). Then $\mathrm{C}^*_{\mathrm{U}\Sigma}(\Gamma) = \bigcap \mathrm{ext}(\mathcal{W})$ *and* $\mathrm{C}^*_{\mathrm{E}\Sigma}(\Gamma) = \bigcup \mathrm{ext}(\mathcal{W})$.

More interestingly, Marek, Treur, and Truszczyński [126] show that there is a close relationship between consistent normal default theories and reasoning from maximal consistent subsets. They also show that the class of consistent prerequisite free normal default theories $\mathcal{P}$ is representationally complete with respect to the class of consistent normal default theories $\mathcal{N}$. Note that $\mathcal{P}$ is a proper subclass of $\mathcal{N}$. Thus the results of Marek, Treur, and Truszczyński can be seen as a strengthening of theorem (3.2.2). The results of Marek, Treur, and Truszczyński can be restated in terms of *representation theory* for default logic. A family of theories $\mathcal{G}$ is said to be trivial if $\Phi \in \mathcal{G}$, else $\mathcal{G}$ is non-trivial. Furthermore $\mathcal{G}$ is said to be *representable* by a default theory $\mathcal{W}$ if $\mathrm{ext}(\mathcal{W}) = \mathcal{G}$.

Theorem 3.2.3
(Marek, Treur, and Truszczyński [126])

1. *If a family of non-trivial theories $\mathcal{G}$ is representable by a default theory $\mathcal{W} \in \mathcal{N}$, then $\mathcal{G}$ is representable by a default theory $\mathcal{W}' \in \mathcal{P}$, i.e. $\mathcal{P}$ is representationally complete with respect to $\mathcal{N}$.*

2. *A family of non-trivial theories $\mathcal{G}$ is representable by a default theory $\mathcal{W} \in \mathcal{N}$ iff there is a set of formulae Γ such that*

$$\mathcal{G} = \{\mathbf{Cn}(\mathcal{A}) : \mathcal{A} \in \mathrm{M}_{\emptyset}(\Gamma)\}$$

Theorem (3.2.3) shows that each consistent normal default theory $\mathcal{W}$ is expressively equivalent to a family of maximal consistent subsets of some set Γ, i.e. the family of closures of these maximal consistent subsets of Γ is precisely the family of extensions of $\mathcal{W}$. But note that Γ need not be unique in general, i.e. it is possible that for some $\Gamma' \neq \Gamma$, $\{\mathbf{Cn}(\mathcal{A}) : \mathcal{A} \in \mathrm{M}_{\emptyset}(\Gamma)\} = \{\mathbf{Cn}(\mathcal{B}) : \mathcal{B} \in \mathrm{M}_{\emptyset}(\Gamma')\}$. The following are immediate corollaries of theorem (3.2.3).

Corollary 3.2.2
1. *For any $\Gamma \subset \Phi$ with $\mathrm{M}_{\emptyset}(\Gamma) \neq \emptyset$, $\mathrm{C}_{\mathrm{E}\emptyset}(\Gamma) = \mathrm{C}^*_{\mathrm{E}\emptyset}(\Gamma) = \bigcup \mathrm{ext}(\mathcal{W})$ and $\mathrm{C}_{\mathrm{U}\emptyset}(\Gamma) = \mathrm{C}^*_{\mathrm{U}\emptyset}(\Gamma) = \bigcap \mathrm{ext}(\mathcal{W})$ for some $\mathcal{W} \in \mathcal{N}$.*

2. *A family of non-trivial theories $\mathcal{G}$ is representable by a default theory $\mathcal{W} \in \mathcal{P}$ iff there is a set of formulae Γ such that*

$$\mathcal{G} = \{\mathbf{Cn}(\mathcal{A}) : \mathcal{A} \in \mathrm{M}_{\emptyset}(\Gamma)\}$$

Proposition 3.2.1

For every $\mathcal{W} \in \mathcal{N}$, $\mathrm{ext}(\mathcal{W})$ can be expressed in two equivalent forms:

$$\{\mathbf{Cn}(\mathcal{A}) : \ \mathcal{A} \in \mathrm{M}_{\emptyset}(\Gamma')\} \quad = \quad \{\mathbf{Cn}(\Sigma \cup \mathcal{B}) : \ \mathcal{B} \in \mathrm{M}_{\Sigma}(\Gamma)\} \tag{3.3}$$

for some consistent Σ and some Γ and Γ'.

Proof:

Consider an arbitrary $\mathcal{W} \in \mathcal{N}$. By (2) of theorem (3.2.3), there exits Γ' such that $\mathrm{ext}(\mathcal{W})$ can be expressed as the LHS of equation (3.3). By (1) of theorem (3.2.3), $\mathcal{P}$ is representationally complete with respect to $\mathcal{N}$. Hence, there exists a $\mathcal{W}' \in \mathcal{P}$ where $\mathcal{W}'$ is extensionally equivalent to $\mathcal{W}$. By theorem (3.2.2), there exits Γ such that $\mathcal{W}'$ can be expressed as the RHS of equation (3.3). ∎

We can give a representational completeness result similar to (1) of theorem (3.2.3). Set $S' = \{(\Gamma, \emptyset) : \ \Gamma \subseteq \Phi\}$ and $S = \{(\Gamma, \Sigma) : \ \Gamma, \Sigma \subseteq \Phi, \ \Sigma \not\vdash \bot\}$. We say that a family of theories $\mathcal{G}$ is representable by a $(\Gamma, \Sigma) \in S$ iff $\mathcal{G} = \{\mathbf{Cn}(\Sigma \cup \mathcal{A}) : \ \mathcal{A} \in \mathrm{M}_{\Sigma}(\Gamma)\}$.

Theorem 3.2.4

If a family of theories $\mathcal{G}$ is representable by some $(\Gamma, \Sigma) \in S$, then $\mathcal{G}$ is representable by some $(\Gamma', \emptyset) \in S'$

Proof:

Let $\mathcal{G}$ be a family of theories representable by some $(\Gamma, \Sigma) \in S$. Then $\mathcal{G}$ is representable by some $\mathcal{W}' \in \mathcal{P}$. Since $\mathcal{P}$ is a subclass of $\mathcal{N}$, we can apply proposition (3.2.1) to $\mathcal{W}'$ and thereby yielding the existence of a Γ' such that $\mathrm{ext}(\mathcal{W}') = \{\mathbf{Cn}(\mathcal{A}) : \mathcal{A} \in \mathrm{M}_{\emptyset}(\Gamma')\}$. This suffices to show that $\mathcal{G}$ is representable by some $(\Gamma', \emptyset) \in S'$. ∎

We can summarise our results with figure (3.1). $\mathcal{P}$, $\mathcal{N}$, S' and S are all expressively equivalent in the sense that a family of theories $\mathcal{G}$ is representable in one of these classes iff $\mathcal{G}$ is representable in all the other classes.

Of course these results show that the formalism of default reasoning is equipped with capabilities for handling inconsistencies. Taking union (intersection) over extensions of a consistent normal default theory corresponds to existential (universal) con-

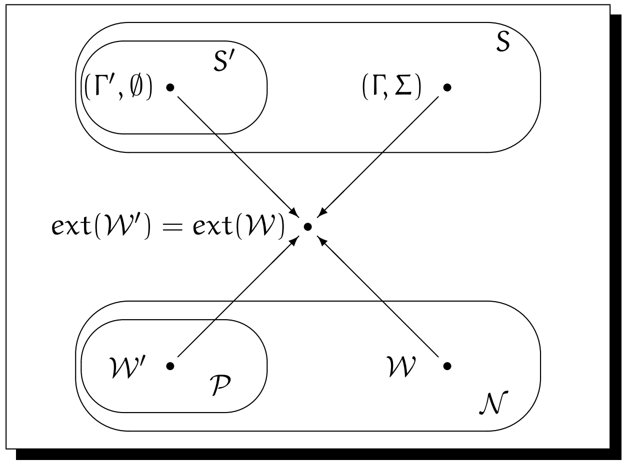

Figure 3.1: Expressive equivalence of $\mathcal{P}$, $\mathcal{N}$, S' and S

sequence in Rescher's sense. But there are also limitations to default reasoning – inconsistencies in the set of facts in a default theory still trivialises the extension. Clearly there is room for improvement here. One possible solution suggested by Hunter in [91] is to replace the underlying classical logic with a weaker paraconsistent logic as the underlying deduction mechanism. Indeed, we can keep most of definition (3.2.1) intact. The only modification we need to make is the replacement of the closure condition in clause 2 of definition (3.2.1) with the deductive closure of a weaker paraconsistent logic. As to which paraconsistent logic should be used, we need not decide *a priori* here. In fact it would seem to be more useful to study and compare the behaviour of the resulting mechanisms obtained by plugging in various paraconsistent logics. We'll leave this analysis and study for future work.

From a preservational point of view, the upshot of theorem (3.2.2) and theorem (3.2.3) is that the kind of analyses and accounts presented in the last chapter has direct counterparts for default theories in $\mathcal{P}$ and $\mathcal{N}$.

3.3 Belnap's Conjunctive Containment

In [8], Belnap considered a strategy to improve the Rescher-mechanism by finding different *articulations* for a set of logical descriptions. Recall that Belnap's main criticism of Rescher is that reasoning with maximal consistent subsets is too syntax dependent on the underlying logical representation. Hence a minor syntactic variant may yield wildly different conclusions. Note that given the results from last section, Belnap's criticism is equally applicable to default reasoning. For the classes $\mathcal{P}$ and $\mathcal{N}$ of default theories, extensions are just closures of maximal consistent subsets. Hence any reasoning strategy involving extensions in these classes is equivalent to reasoning with maximal consistent subsets.

Belnap's main idea is that given a set of input premises Γ we can *pre-process* Γ with certain closure operations so that the content of the input premises can be made explicit and information not involved in any inconsistency can be isolated. Once this is done, we can then apply the Rescher-mechanism to reason with the extended set. Indeed Belnap's suggestion is not fundamentally different from the methodology of *knowledge compilation* in AI (see [106; 107]). In knowledge compilation, the general aim is to give a sound and complete *translation* of information represented in a general language to a sub-language with lower complexity. The translation is done *off-line* so that the computational cost of inference is shifted from run time query-answering to off-line compilation. Thus in knowledge compilation, reasoning can be viewed as a two stage process involving both data preparation and formal deduction from prepared data. For Belnap however, the concern is not so much to reduce the computational cost but to reduce the effect of syntactic variations on inferences from the innocent part of the information. Clearly this can be seen as a form of data preparation. It is instructive to recap Belnap's reasons for rejecting the use of the first degree entailment (FDE) of the relevant logic R, Parry's analytic implication (AI) and Angell's analytic containment (AC) as candidate closures for the input premises:

Example 3.3.1
$\Gamma = \{p,\ \neg p,\ q\}$

In FDE and Parry's AI we have, respectively

$$\vdash_{\mathsf{FDE}} A \rightarrow A \vee B \qquad\qquad \vdash_{\mathsf{AI}} A \wedge B \rightarrow A \vee \neg B$$

Hence the closure of Γ under either FDE or analytic implication yields $\neg p \vee \neg q$ which conspires together with p to prevent q from being derived as a U-consequence.

Example 3.3.2
$\Gamma = \{p,\ \neg p,\ q,\ r \vee \neg q\}$

In Angell's AC we have

$$\vdash_{\mathsf{AC}} A \wedge (B \vee C) \leftrightarrow A \wedge (B \vee C) \wedge (A \vee C)$$

Hence from Γ we get $\vdash_{AC} \neg p \wedge (r \vee \neg q) \rightarrow \neg p \vee \neg q$. Once again $\neg p \vee \neg q$ conspires together with p to prevent q from being derived as a U-consequence.

In each of these cases, the use of a certain version of disjunction introduction results in the introduction of additional inconsistencies. Indeed this is symptomatic of the kind of difficulties involved in the use of disjunction introduction in the presence of inconsistencies. To overcome this problem, Belnap proposes the use of *conjunctive containment*. First we have the following definitions. To simplify the matter, we'll assume that our language is restricted to the truth functional connectives $\{\neg, \wedge, \vee\}$.

Definition 3.3.1
A subformula B *of a given formula* A *is said to be an even subformula if it is within the scope of zero or an even number of negations, otherwise it is said to be odd.*

Definition 3.3.2
Belnap's replacement rules are given as follows:

$$[*] \; \frac{\ldots (B \wedge C) \ldots}{\ldots B \ldots \qquad \ldots C \ldots}$$

provided that $(B \wedge C)$ *is an even subformula.*

$$[\#] \; \frac{\ldots (B \vee C) \ldots}{\ldots B \ldots \qquad \ldots C \ldots}$$

provided that $(B \vee C)$ *is an odd subformula.*

Clearly for any given A we can built a finite binary tree $\mathcal{T}$ such that

1. the root of $\mathcal{T}$ is just A,
2. each branching is an application of either [#] or [$*$],
3. a node is either the root of $\mathcal{T}$ or a formula obtained by [#] or [$*$], and
4. the leaves or end points are formulae which contain no even subformulae of the form $(B \wedge C)$ and no odd subformulae of the form $(B \vee C)$

For convenience we shall draw a tree with the root at the bottom and all branches extending upward, i.e. we apply the replacement rules as if they are upside down. Since the order in which we apply [#] and [$*$] can be permuted, clearly such a tree is not unique for a given A in general. But there can be at most finitely many such trees for a given A. Thus we can associate with each A a finite set of trees $\{\mathcal{T}_1^A, \ldots, \mathcal{T}_n^A\}$

where each $\mathcal{T}_i^A$ is a finite binary tree built in the prescribed way. We'll call these the Belnap trees associated with A.

Lemma 3.3.1
Let $(B \wedge C)$ *be a zero subformula in* $D = \ldots (B \wedge C) \ldots$, *then* D *is classically equivalent to* $D' = ((\ldots B \ldots) \wedge (\ldots C \ldots))$.

Proof:
Since $(B \wedge C)$ is a zero subformula, we can equivalently transform D into the following form:

$$D_0 = \overbrace{(\ldots(}^{n+1} B \wedge C) *_1 A_1) *_2 A_2) *_3 \ldots) *_n A_n)$$

where $*_i$ is either $\wedge$ or $\vee$.

Since both $\wedge$ and $\vee$ are commutative, we note any step in the transformation from D to D_0 is reversible and equivalence preserving. We'll denote the transformation from D to D_0 as T and the reverse of T as T′. We'll show by induction on the depth d of $(B \wedge C)$, defined in terms of the number of '(' to the left of $B \wedge C$, that D_0 is equivalent to

$$D_0' = [\overbrace{(\ldots(}^{n} B *_1 A_1) *_2 \ldots) *_n A_n) \wedge \overbrace{(\ldots(}^{n} C *_1 A_1) *_2 \ldots) *_n A_n)]$$

For the basis $d = 1$: this is trivial since $(B \wedge C)$ is equivalent to itself. For the inductive step, We'll make the assumption that the statement holds for $d = k$ and show that it holds for the case when $d = k + 1$. Since we assume that $d = k + 1$, D_0 must be of the form:

$$\overbrace{(\ldots(}^{k+1} B \wedge C) *_1 A_1) *_2 A_2) *_3 \ldots) *_k A_k)$$

By the induction hypothesis, the following subformula of D_0

$$\overbrace{(\ldots(}^{k} B \wedge C) *_1 A_1) *_2 A_2) *_3 \ldots) *_{k-1} A_{k-1})$$

is equivalent to

$$[\overbrace{(\ldots(}^{k} B *_1 A_1) *_2 \ldots) *_{k-1} A_{k-1}) \wedge \overbrace{(\ldots(}^{k} C *_1 A_1) *_2 \ldots) *_{k-1} A_{k-1})]$$

Hence D_0 must be equivalent to

$$E = \Big([\, \overbrace{(\ldots(}^{k} B *_1 A_1) *_2 \ldots) *_{k-1} A_{k-1}) \wedge \overbrace{(\ldots(}^{k-1} C *_1 A_1) *_2 \ldots) *_{k-1} A_{k-1})\,] *_k A_k \Big)$$

There are two cases to consider: either $*_k$ is $\wedge$ or $\vee$. In the first case E is equivalent to

$$[\, \overbrace{(\ldots(}^{k+1} B *_1 A_1) \ldots *_{k-1} A_{k-1}) \wedge A_k) \wedge \overbrace{(\ldots(}^{k+1} C *_1 A_1) \ldots *_{k-1} A_{k-1}) \wedge A_k)]$$

In the later case, using distribution of $\vee$ over $\wedge$, E is equivalent to

$$[\, \overbrace{(\ldots(}^{k+1} B *_1 A_1) \ldots *_{k-1} A_{k-1}) \vee A_k) \wedge \overbrace{(\ldots(}^{k+1} C *_1 A_1) \ldots *_{k-1} A_{k-1}) \vee A_k)]$$

This suffices to show that D_0 and D'_0 are equivalent. To complete the proof we make use of the fact that the transformation T from D to D_0 is reversible. Hence by applying the reverse transformation T' to the left and right conjuncts of D'_0, the equivalence of D and D' follows. ∎

Proposition 3.3.1
If E and F are obtained from A by an application of either $[*]$ or $[\#]$, then $E \wedge F$ is classically equivalent to A.

Proof:
If E and F are obtained from A by an application of $[*]$, then E and F must be obtained via an even subformula $(B \wedge C)$ of A. Since $(B \wedge C)$ is even, repeat applications of pushing negations onto $(B \wedge C)$ will result in an even occurrences of negation in front of $(B \wedge C)$. By double negation elimination we can transform A into an equivalent formula A' where $(B \wedge C)$ is a zero subformula. Using lemma (3.3.1) and reversing the relevant transformation steps, the desired result follows.

If E and F are obtained via $[\#]$, then E and F must be obtained via an odd subformula $(B \vee C)$ of A. Since $(B \vee C)$ is odd, repeat application of pushing negation onto $(B \vee C)$ will result in an odd occurrences of negation in front of $(B \vee C)$. Using double negation elimination repeatedly and pushing the remaining negation into $(B \vee C)$ will result in an even subformula $(\neg B \wedge \neg C)$. Again applying lemma 3.3.1 and reversing all relevant transformation steps, the desired result follows. ∎

Corollary 3.3.1

For any Belnap tree $\mathcal{T}^A$ associated with A, A is classically equivalent to the conjunction of all the leaves in $\mathcal{T}^A$.

Definition 3.3.3

Belnap's Closure, C_B, on a given A is defined as follows: $D \in C_B(A)$ iff D is a node of some Belnap tree associated with A. For a given set of formulae Γ, $C_B(\Gamma) = \{D \in C_B(A) : A \in \Gamma\}$.

A set Γ conjunctively contains A in the strict sense iff $A \in C_B(\Gamma)$, i.e. A is a node of some Belnap tree associated with some $A \in \Gamma$.

The extended Belnap's Closure, C_B^+ on a set Γ is defined as follows: $A \in C_B^+(\Gamma)$ iff every member of $C_B(A)$ is classically equivalent to a conjunction of some members of $C_B(\Gamma)$.

Alternatively we may define $C_B(\Gamma)$ simply as the least superset of Γ that is closed under [$*$] and [#]. The use of Belnap's trees gives us an easy way to visualise the underlying mechanism: each piece of information A is *conjunctively eliminated* at each level of a Belnap's tree until all hidden conjunctions are eliminated; since [$*$] and [#] preserve the model of their premises, all information implicitly encoded in A is successively passed on to the next level in its Belnap's tree.

With respect to the extended Belnap's closure C_B^+, the basic idea is to regain some limited form of conjunction introduction with members of C_B while adding all those that are classically equivalent to these conjunctions without creating unexpected nonequivalence.

Example 3.3.3

Let $\Gamma = \{p\}$. We have $p \wedge p \in C_B^+(\Gamma)$ but $p \vee (p \wedge q) \notin C_B^+(\Gamma)$ even though $p \wedge p$ is classically equivalent to $p \vee (p \wedge q)$. Note that although we have $p \vee q \in C_B(\{p \vee (p \wedge q)\})$, $p \vee q$ is not classically equivalent to any conjunction of members of $C_B(\Gamma)$.

Fact 3.3.1

C_B and C_B^+ are closure operators in the sense of Tarski, i.e. they satisfy inclusion, monotonicity and idempotence. C_B^+ is an extension of C_B, i.e. for any $\Gamma \subseteq \Phi$, $C_B(\Gamma) \subseteq C_B^+(\Gamma)$. Moreover they distribute over union, i.e. $C_B(\Gamma \cup \Gamma') = C_B(\Gamma) \cup C_B(\Gamma')$

While we may think of Belnap's closure C_B as an articulation of a set of formulae, the extension C_B^+ is a *proper* E-*equivalent extension* of C_B in the following sense:

Definition 3.3.4
A closure operator C is a proper E-equivalent extension of a closure operator C' iff for any premise set $\Gamma \subseteq \Phi$,

1. *$C(\Gamma)$ and $C'(\Gamma)$ have exactly the same set of existential-consequences, i.e. $C^*_E(C(\Gamma)) = C^*_E(C'(\Gamma))$.*

2. *$C'(\Gamma) \subseteq C(\Gamma)$*

3. *$C(C'(\Gamma)) = C(\Gamma)$*

If condition (1) holds for C, then we say that C is E-equivalent to C'.

Note that in our definition, we have dropped the reference to any constraint set Σ. By the representational completeness of theorem (3.2.4), there is no loss of generality here. By modifying clause (1) of our definition, we can obtain the corresponding notions of proper U, A, S, L-equivalent extensions of a given closure operator. More generally for $x \in \{E, U, A, S, L\}$, two sets Γ and Γ' are x-equivalent if Γ and Γ' have exactly the same set of x-consequences. The following lemma can be used to show that C^+_B is a proper E and U-equivalent extension of C_B:

Lemma 3.3.2
For an arbitrary but fixed Γ, let $M(C_B(\Gamma))$ and $M(C^+_B(\Gamma))$ be the collections of maximal consistent subsets of $C_B(\Gamma)$ and $C^+_B(\Gamma)$ respectively. Then there is a bijection f with domain $M(C_B(\Gamma))$ and range $M(C^+_B(\Gamma))$ such that for any $\mathcal{A} \in M(C_B(\Gamma))$, $f(\mathcal{A})$ is classically equivalent to $\mathcal{A}$.

Proof:
Let Γ be arbitrary but fixed. Let

$$\begin{aligned} M(C_B(\Gamma)) &= \{\mathcal{A}_i : i \in I\} \\ M(C^+_B(\Gamma)) &= \{\mathcal{B}_j : j \in J\} \end{aligned}$$

We observe that

1. for each $i \in I$ there exists a $j \in J$ such that $\mathcal{A}_i \subseteq \mathcal{B}_j$: by the consistency of $\mathcal{A}_i$ and the fact that $C_B(\Gamma) \subseteq C^+_B(\Gamma)$.

2. for each $i \in I$ there is exactly one $j \in J$ such that $\mathcal{A}_i \subseteq \mathcal{B}_j$: from (1) the existence of such a $\mathcal{B}_j$ is guaranteed for each arbitrary but fixed $i \in I$. Toward a contradiction

assume that for some $k \in J$, $k \neq j$, $\mathcal{A}_i \subseteq \mathcal{B}_k$. Note that since $\mathcal{B}_j \neq \mathcal{B}_k$, $\mathcal{B}_j \cup \mathcal{B}_k$ is inconsistent. Hence there exists $D_1, \ldots, D_m \in \mathcal{B}_j$ and $E_1, \ldots, E_n \in \mathcal{B}_k$ such that

$$D_1 \wedge \ldots \wedge D_m \vdash \neg(E_1 \wedge \ldots \wedge E_n)$$

We claim that every member of $C_B(D_1) \cup \ldots \cup C_B(D_m)$ must be classically equivalent to a conjunction of some members of $\mathcal{A}_i$. Suppose not. Then there must be a member of $C_B(D_1) \cup \ldots \cup C_B(D_m)$ classically equivalent to a conjunction of members involving elements of $(C_B(\Gamma) \setminus \mathcal{A}_i)$. But this is impossible since by the maximal consistency of $\mathcal{A}_i$ any $A \in C_B(\Gamma) \setminus \mathcal{A}_i$ is inconsistent with $\mathcal{A}_i$ and this would imply that $\mathcal{A}_i$ is inconsistent with $\mathcal{B}_j$. Similar argument also shows that every member of $C_B(E_1) \cup \ldots \cup C_B(E_n)$ must be classically equivalent to a conjunction of some members of $\mathcal{A}_i$. But this clearly contradicts the consistency of $\mathcal{A}$. Hence $\mathcal{B}_j$ cannot be distinct from $\mathcal{B}_k$ afterall.

3. for no $i, i' \in I$, $i \neq i'$ do we have $\mathcal{A}_i \subseteq \mathcal{B}_j$ and $\mathcal{A}_{i'} \subseteq \mathcal{B}_j$ for some $j \in J$: by the consistency of each $\mathcal{B}_j$.

4. for each $j \in J$ there exists an $i \in I$ such that $\mathcal{A}_i \subseteq \mathcal{B}_j$: it is straightforward to verify that for each $j \in J$, $\mathcal{B}_j \cap C_B(\Gamma)$ is a maximal consistent subset of $C_B(\Gamma)$.

We now define the function $f : C_B(\Gamma) \longrightarrow C_B^+(\Gamma)$ as follows: for each $i \in I$

$$f(\mathcal{A}_i) = \mathcal{B}_j \Leftrightarrow \mathcal{A}_i \subseteq \mathcal{B}_j$$

for some $j \in J$. Clearly by observation (2), f is a well defined function. By observation (4), f is surjective. By observation (3) f is injective. Hence f is a bijection.

Finally to show that for every $i \in I$, $\mathcal{A}_i$ and $f(\mathcal{A}_i)$ are classically equivalent, it suffices to observe that the argument for observation (2) establishes that for every $A \in f(\mathcal{A}_i)$, every member of $C_B(A)$ is classically equivalent to a conjunction of some members of $\mathcal{A}_i$. ∎

Theorem 3.3.1
C_B^+ *is a proper* E *and* U*-equivalent extension of* C_B.

Proof:
By lemma (3.3.2), condition (1) of definition (3.3.4) is clearly satisfied. Moreover $C_B(\Gamma) \subseteq C_B^+(\Gamma)$ clearly holds. It remains to verify that for any Γ, $C_B^+(C_B(\Gamma)) = C_B^+(\Gamma)$

($\supseteq$): Since C_B and C_B^+ are both Tarskian closure operators, we have $\Gamma \subseteq C_B(\Gamma)$ and hence $C_B^+(\Gamma) \subseteq C_B^+(C_B(\Gamma))$.

($\subseteq$): If $A \in C_B^+(C_B(\Gamma))$, then for every $B \in C_B(A)$, there are $C_1, \ldots, C_n \in C_B(C_B(\Gamma))$ such that B is classically equivalent to $C_1 \wedge \ldots \wedge C_n$. But $C_B(C_B(\Gamma)) = C_B(\Gamma)$, hence $A \in C_B^+(\Gamma)$ as required. ∎

Corollary 3.3.2
For any Γ and consistent $\Sigma \subseteq \Phi$, let $\mathcal{D}$ and $\mathcal{D}'$ be default theories defined as follows:

$$\mathcal{D} = \langle\{\frac{:A}{A} : A \in C_B(\Gamma)\}, \Sigma\rangle \qquad \mathcal{D}' = \langle\{\frac{:B}{B} : B \in C_B^+(\Gamma)\}, \Sigma\rangle$$

Then $\mathrm{ext}(\mathcal{D}) = \mathrm{ext}(\mathcal{D}')$

The notion of E-equivalence is an important idea and by theorem (3.2.3) it is related to the notion of extension equivalence between default theories. Lemma (3.3.2) clearly gives a sufficient condition for E-equivalence – two sets of formulae (with arbitrary cardinalities) are E-equivalent if a bijection of the suitable sort exists between the two collections of maximal consistent subsets of the two sets. But it is unclear that this is also necessary in cases where infinite cardinalities are considered. In [156], Rescher and Manor give the necessary and sufficient conditions for E-equivalence for *finitely generated* sets of formulae. But no general characterisation is given there. For sets with finitely many maximal consistent subsets (though not necessarily finitely generated in Rescher's sense), the following proposition gives a necessary condition for their E-equivalence:

Proposition 3.3.2
Let $|M(\Gamma)| < \omega$. If Γ' is E-equivalent to Γ, then $|M(\Gamma')| = |M(\Gamma)|$. Equivalently, for any Γ and Γ', if $|M(\Gamma')|$ and $|M(\Gamma)|$ are finite and $|M(\Gamma')| \neq |M(\Gamma)|$, then Γ and Γ' are not E-equivalent.

Proof:
Without loss of generality, we may assume that Γ and Γ' are such that $|M(\Gamma)| = n$ and $|M(\Gamma')| = m$ where $m < n$. Towards a contradiction, we assume that Γ and Γ' are E-equivalent. We let $M(\Gamma) = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ and $M(\Gamma') = \{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$. Since members of $M(\Gamma')$ are pairwise inconsistent, by the standard compactness theorem, there exist formulae $A_1, \ldots, A_n$ such that for every $i \leq n$,

1. $\mathcal{A}_i \vdash A_i$ and

2. $A_i \vdash \bigwedge_{i \neq j} \neg A_j$

By our reductio assumption Γ and Γ' are E-equivalent and hence by (1) above for each $i \leq n$ there exists a $k \leq m$ such that $\mathcal{B}_k \vdash A_i$. However by our initial assumption $m < n$ and thus by the pigeonhole principle, there exists a $t \leq m$ such that $\mathcal{B}_t \vdash \bigwedge_{i \leq n} \neg A_i$. Clearly by the consistency of each $\mathcal{A}_i \in M(\Gamma)$, $\mathcal{A}_i \nvdash \bigwedge_{i \leq n} \neg A_i$. But this contradicts the assumption that Γ and Γ' are E-equivalent. This suffices to show that Γ and Γ' are not E-equivalent on the assumption that $n \neq m$. ∎

Thus for any two finite sets, we can test for their non-E-equivalence by simply counting their number of maximal consistent subsets. In previous chapter, such a counting function λ was introduced and studied. It was further argued that the λ value of a set of formulae may be used as a way to measure the amount of inconsistency in the set. Intuitively, if the λ values of two finite sets (with the same cardinality) differ then the amounts of inconsistency in these two sets also differ. This intuition is justified by the fact that if $\lambda(\Gamma) = k \neq \omega$, then by taking the union of each distinct pair of maximal consistent subsets there are at least $\frac{k(k-1)}{2}$ many ways of generating inconsistent subsets of Γ. As such the λ function may be useful as a tool for analyzing inconsistent data.

3.3.1 Maximal Equivalent Extension

In [8] Belnap considered the interesting possibility of finding a strongest, i.e. maximal, closure operator C that U-equivalently extends C_B. In particular, Belnap proposed C_B^+ as a candidate. The following shows that C_B^+ fails to be such a maximal closure operator.

Proposition 3.3.3
C_B^+ *is not a maximal closure that* U*-equivalently extends* C_B*.*

Proof:
It suffices to find a closure operator that properly extends C_B^+ and U-equivalently extends C_B. For any Γ let,

$$C^*(\Gamma) = C_B^+(\Gamma) \cup \top$$

We claim that $C_B^+(\Gamma) \cup \top = C_B^+(\Gamma \cup \top)$:

($\subseteq$): trivial since C_B^+ is a closure operator.

($\supseteq$): let $A \in C_B^+(\Gamma \cup \top)$. Then each $B \in C_B(A)$ is classically equivalent to the conjunction of some $C_1, \ldots, C_n \in C_B(\Gamma \cup \top)$. But $C_B(\Gamma \cup \top) = C_B(\Gamma) \cup \top$ by the distributivity of C_B, so we have 3 cases to consider:

1. $C_1, \ldots, C_n \in C_B(\Gamma)$. Then $A \in C_B^+(\Gamma)$ and hence $A \in C_B^+(\Gamma) \cup \top$.
2. $C_1, \ldots, C_n \in \top$. Then A is a tautology and hence $A \in C_B^+(\Gamma) \cup \top$.
3. $C_1, \ldots, C_i \in C_B(\Gamma)$ and $C_{i+1} \ldots C_n \in \top$. Then clearly A is equivalent to $C_1 \wedge \ldots \wedge C_i$. Hence $A \in C_B^+(\Gamma)$.

By our claim C^* is a closure operator. We verify that $C^*(C_B(\Gamma)) = C^*(\Gamma)$:

($\supseteq$): since $\Gamma \subseteq C_B(\Gamma)$, $C^*(\Gamma) \subseteq C^*(C_B(\Gamma))$ holds as required.

($\subseteq$): by the usual closure properties of C_B, we have $C_B(C_B(\Gamma) \cup \top) = C_B(\Gamma \cup \top)$. Hence $A \in C_B^+(C_B(\Gamma) \cup \top)$ implies that $A \in C_B^+(\Gamma \cup \top)$. Hence $C^*(C_B(\Gamma)) \subseteq C^*(\Gamma)$.

It remains to show that $C_B(\Gamma)$ and $C^*(\Gamma)$ have the same set of U-consequences. But this is trivial since tautologies are trivial consequences of any subset. Finally to see that C^* properly extends C_B^+, we note that $C_B^+(\Gamma) \subset C^*(\Gamma)$ holds for any Γ such that $C_B(\Gamma)$ contains no tautology. ∎

Though technically correct, our proposition is unremarkable since it is straightforward to show that C^* properly extends C_B^+ and U-equivalently extends C_B. However, we may think that we can continue the trick by adding every false sentence to Γ and then closing the resulting set under C_B^+. But this is clearly not a U-equivalent extension of C_B. If we begin with a consistent set, e.g. $\{p\}$ and then add the false sentence $p \wedge \neg p$, we can no long obtain p as a U-consequence.

Another minor observation is that if there are countably many propositional atoms, then there are at least countably many trivial U-equivalent extensions of C_B between C_B and C^*: for each p_i we can simply add $p_i \vee \neg p_i$ to Γ and then close it under C_B to obtain a U-equivalent extension.

3.4 An Improvement to Belnap's Strategy

One of the main motivations for Belnap to introduce C_B is to provide a standard way to isolate the effect of the inconsistencies in a set. Recall that the rules [$*$] and [#] are replacement rules of the form:

$$[*]\ \frac{\ldots(B \wedge C)\ldots}{\ldots B \ldots \qquad \ldots C \ldots} \qquad\qquad [\#]\ \frac{\ldots(B \vee C)\ldots}{\ldots B \ldots \qquad \ldots C \ldots}$$

where the $[*]$ rule applies if $(B \wedge C)$ is even and the [#] rule applies if $(B \vee C)$ is odd. These rules are introduced by Belnap specifically to eliminate concealed conjunctions. In [87], Horty explicitly endorsed a similar strategy for handling inconsistent instructions using modalized versions of these replacement rules:

$$[\Box *]\ \frac{\Box(\ldots(B \wedge C)\ldots)}{\Box(\ldots B \ldots) \qquad \Box(\ldots C \ldots)} \qquad\qquad [\Box \#]\ \frac{\Box(\ldots(B \vee C)\ldots)}{\Box(\ldots B \ldots) \qquad \Box(\ldots C \ldots)}$$

where again the even and odd restrictions apply to the respective rule.

To appreciate the significance of these rules, we quote a remark of Belnap:

> Since different ways of articulating our beliefs . . . give different results under Rescher's proposal and since we do not want this, evidently we have to have some views about which articulations we most want to reflect . . . Policy: try to reflect *maximum* articulation. . . . if we maximally articulate . . . we may be able to isolate the effect of its contradiction, . . . [o]r, which seems just as important, we may be able to block a consequence by freeing for use some conjunct of a conjunction which is itself not consistently available . . . (page 545 [8])

In light of the $[*]$ and [#] rules, *maximum* articulation here is cashed out in terms of conjunction elimination. In certain cases, this seems to be just the right remedy. Consider for instance:

Example 3.4.1

Let $A = (p \wedge \neg p) \wedge (q \vee r)$ and let $\Gamma = \{A,\ \neg r,\ \neg p\}$

Applying the $[*]$ rule to A we get

$$\dfrac{\dfrac{p \qquad \neg p}{p \wedge \neg p} \qquad q \vee r}{(p \wedge \neg p) \wedge (q \vee r)}$$

Figure 3.2: Belnap tree for $(p \wedge \neg p) \wedge (q \vee r)$

In example (3.4.1), all conjunctions are eliminated to maximally articulate the information encoded by A. The result is that the inconsistency with respect to p would neither interfere with $\neg r$ nor $q \vee r$. We note however that an imprudent use of $[*]$ may result in duplication and thereby increase the size of the tree.

$$\dfrac{\dfrac{p \qquad q \vee r}{p \wedge (q \vee r)} \qquad \dfrac{\neg p \qquad q \vee r}{\neg p \wedge (q \vee r)}}{(p \wedge \neg p) \wedge (q \vee r)}$$

Figure 3.3: Belnap tree for $(p \wedge \neg p) \wedge (q \vee r)$

However, our main concern here is not with efficiency. Our main concern is that there are cases in which $[*]$ and [#] cannot eliminate conjunctions without a *detour* in using the distributive properties of $\wedge$ over $\vee$ and vice versa. Consider for instance a slight variant of example (3.4.1):

Example 3.4.2
Let $B = \neg[(p \vee \neg p) \vee \neg q] \vee \neg[(p \vee \neg p) \vee \neg r)]$ and $\Gamma' = \{B, \neg r, \neg p\}$

Assuming that we have the usual double negation elimination rule, contraction for $\vee$ and $\wedge$, and commutative and associative rules for $\vee$ and $\wedge$, we can now apply [#] to B in example (3.4.2) to obtain the Belnap tree for B (figure (3.4).

$$\dfrac{\dfrac{\dfrac{p \qquad \neg p}{\neg(p \vee \neg p)} \qquad \dfrac{\neg p \vee r \qquad p \vee r}{\neg[(p \vee \neg p) \vee \neg\neg r]}}{\neg(p \vee \neg p) \vee \neg[(p \vee \neg p) \vee \neg r]} \qquad \dfrac{\dfrac{q \vee r}{\neg\neg q \vee \neg\neg r} \qquad \dfrac{p \vee q \qquad \neg p \vee q}{\neg\neg q \vee \neg(p \vee \neg p)}}{\neg\neg q \vee \neg[(p \vee \neg p) \vee \neg r]}}{\neg[(p \vee \neg p) \vee \neg q] \vee \neg[(p \vee \neg p) \vee \neg r]}$$

Figure 3.4: Belnap tree for $\neg[(p \vee \neg p) \vee \neg q] \vee \neg[(p \vee \neg p) \vee \neg r)]$

This is not entirely satisfactory because in taking an unnecessary detour, we have produced additional disjunctive information. Belnap's initial objection against the use of relevant implication and analytic implication is precisely that closures under these implications are too liberal in generating disjunctive information. The point is that disjunctive information can interact with inconsistencies in such a way that further inconsistencies can be produced. In the presence of $\neg r$ for instance, $p \vee r$ and $\neg p$ form an inconsistent triad. Comparing this with figure (3.2) however, $\neg r$ remains innocent. We note further that the *distributivity* of $\wedge$ over $\vee$ and $\vee$ over $\wedge$ are built into Belnap's replacement rules – we cannot avoid the use of distributivity with these rules.

While our previous example demonstrates how distributivity is used in the context of implicit conjunction, our next example shows that the same is true of explicit conjunction:

Example 3.4.3
Let $C = (p \wedge q) \vee (p \wedge r)$ and $\Gamma = \{C, \neg r, \neg p\}$

Applying the [$*$] rule to C we get the following Belnap tree:

$$\dfrac{\dfrac{\mathbf{p} \qquad \mathbf{p \vee r}}{p \vee (p \wedge r)} \qquad \dfrac{\mathbf{p \vee q} \qquad q \vee r}{q \vee (p \wedge r)}}{(p \wedge q) \vee (p \wedge r)}$$

Figure 3.5: Belnap tree for $(p \wedge q) \vee (p \wedge r)$

$$\dfrac{\dfrac{p \qquad q \vee r}{p \wedge (q \vee r)}}{(p \wedge q) \vee (p \wedge r)}$$

Figure 3.6: Factoring for $(p \wedge q) \vee (p \wedge r)$

The example here is similar to the previous case: $p \vee r$ conspires with $\neg p$ in Γ to prevent $\neg r$ from being derived as an U-consequence. Contrasting this with the case where we use a more direct route to eliminate conjunction while keeping disjunctive information to a minimal, we get a very different result. But what can we say coherently about these examples? There are two possible options:

1. our examples do encode different information in each instance and hence conjunctive containment merely makes explicit the difference. This is reflected in the production of different disjunctive information under conjunctive containment.

2. our examples do not encode substantively different information in each instance. The fact that their conjunctive closures differ shows that conjunctive containment still over generates – in particular it over generates by producing too much disjunctive information.

Our intuition in this matter may not run very deep. Indeed there may not be any definitive reason to settle for one over the other. While option (1) is certainly a coherent position (and we suspect this is the option Belnap is likely to take), we would like to explore option (2) here and flesh out an account where examples (3.4.1), (3.4.2), and (3.4.3) do not yield different U-consequences while their maximal articulations are narrower then conjunctive containment.

3.4.1 Logic Minimisation

We start with example (3.4.3) first. We note that the set of leaves $\Delta' = \{p,\ q \vee r\}$ in figure (3.6) is consistent and classically equivalent to the set of leaves $\Delta = \{p,\ p \vee r,\ p \vee q,\ q \vee r\}$ in figure (3.5). However, they clearly differ in the way in which they interact with the remaining members of Γ. Δ would generate more inconsistent subsets when added to Γ then would Δ'. Furthermore, we note that every member of Δ' is a *prime implicate* of C. We'll briefly recap some of the standard definitions:

Definition 3.4.1
A literal is either a propositional atom or the negation of a propositional atom. A disjunction of literals is said to be a clause. A clause D *is an implicate of* A *iff* $A \models D$*. An implicate* D *of* A *is prime iff for all implicates* D' *of* A *if* $D' \models D$ *then* $D \models D'$*. A set of prime implicates* $\{D_1, \ldots, D_n\}$ *of* A *is complete iff* $\{D_1, \ldots, D_n\} \models A$*. A set of prime implicates* $\{D_1, \ldots, D_n\}$ *is independent iff for no* D_i *do we have* $\{D_1 \ldots, D_{i-1}, D_{i+1} \ldots D_n\} \models D_i$

The notion of prime implicate was introduced by Quine in [142; 143; 144]. Quine was interested in simplifying truth functions and he showed that notion of prime implicate plays a central role in simplifying truth functions and thereby contributed directly to the minimisation and design of digital circuits. The emphasis on minimisation stems from the days when the production of logic gates was expensive and required considerable physical space and power. With the advent of semiconductor processes and VSLI, it is of course no longer a central concern to reduce the actual gate count for a system. Circuit design today is more concerned with physical space allocation, reliability and the correctness of a system. Interest in the use of prime implicates in circuit design has decreased considerably as a result. But the notion of prime implicate enjoys a renewed interest in recent years in light of works by de Kleer *et al* in logic based diagnostic systems [56; 57; 58].

Returning to our example however, it is easy to see that Δ' is a complete and independent set of prime implicates of C. It is thus natural to take Δ' to be the maximal articulation of C in example (3.4.3). Our choice can be justified on the grounds that

- Δ' is a more compact representation of C,

- Δ' minimises redundancies and disjunctive information, and
- Δ' minimises interference with $\neg r$

However, the standard notion of prime implicate would not be able to handle example (3.4.2) since definition (3.4.1) uses a classical notion of consequence and thus inconsistent formulae would have the same (complete equivalence class) of prime implicates – namely the empty clause $\emptyset$. But this is exactly what we are trying to avoid in the first place. However, there is a straightforward way to amend the situation – we use a *relevant* notion of prime implicate:

Definition 3.4.2
A clause D *is a relevant implicate of* A *iff* $A \models_{FDE} D$*. A relevant implicate* D *of* A *is prime iff for all relevant implicates* D' *of* A *if* $D' \models_{FDE} D$ *then* $D \models_{FDE} D'$*. A set of relevant prime implicates* $\{D_1, \ldots, D_n\}$ *of* A *is complete iff* $D_1 \wedge \ldots \wedge D_n \models_{FDE} A$*. A set of relevant prime implicates* $\{D_1, \ldots, D_n\}$ *is independent iff for no* D_i *do we have* $D_1 \wedge \ldots \wedge D_{i-1} \wedge D_{i+1} \wedge \ldots \wedge D_n \models_{FDE} D_i$*. We say that two sets of formulae* Γ *and* Δ *are* FDE *equivalent, written as* $\Gamma \equiv_{FDE} \Delta$*, iff* $\bigwedge \Gamma \models_{FDE} \bigwedge \Delta$ *and* $\bigwedge \Delta \models_{FDE} \bigwedge \Gamma$*.*

The dual notion of *relevant prime implicant* of a given formula A can easily be defined: a relevant implicant of a formula A is a *cube* C (conjunction of literals) which FDE-entails A. In addition, C is *prime* if it is a minimal cube that FDE-entails A.

Since FDE is a paraconsistent logic, the empty clause is not a FDE-consequence of any (non-empty) inconsistent formula.

Proposition 3.4.1
For no A *do we have* $A \models_{FDE} \emptyset$*.*

Proof:
Using the standard (4-valued) ambi-valuation of Dunn in [66] (also see appendix (A) for more details), we can verify the existence of a 4-valued assignment v with $1 \notin v(\emptyset)$ while $1 \in v(A)$ for any $A \neq \emptyset$ ∎

Given that *resolution* is not a valid form of inference in relevant logic in general, it is easy to see that the set of classical prime implicates (PI) and the set of relevant prime implicates (RPI) may be distinct for a given formula:

Example 3.4.4
$A = (q \vee r) \wedge ((p \vee q) \wedge (\neg p \vee q))$

	RPI	PI
$q \vee r$	$\surd$	$\times$
q	$\times$	$\surd$

Figure 3.7: RPI and PI of A

Although not every RPI of a given A is a PI of A, it is easy to see that:

Proposition 3.4.2
Every RPI *of a given* A *is a classical implicate of* A*.*

Proof:
It suffices to observe that $\models_{FDE} \subset \models$. ∎

Since any two complete independent sets of relevant prime implicates of a given formula must be FDE equivalent, we can treat them as unique up to equivalence. We'll use the notation $RPI(A)$ to denote any such complete independent set of relevant prime implicates of A. Similarly we use $PI(A)$ for the complete independent set of classical prime implicates of A. We note that $RPI(A)$ is a minimal set (ordered under $\subseteq$) that is both complete and independent. In classical logic, two formulae are equivalent iff their prime implicates are equivalent. This is also true with respect to FDE formulae:

Proposition 3.4.3
For any A *and* B*,* $A \models_{FDE} B$ *and* $B \models_{FDE} A$ *iff* $RPI(A) \equiv_{FDE} RPI(B)$.

Proof:
($\Rightarrow$): Suppose A and B are FDE equivalent. Let $RPI(A) = \{D_1, \ldots, D_m\}$ and $RPI(B) = \{E_1, \ldots, E_n\}$. By the transitivity of $\models_{FDE}$ we have, $\bigwedge_{i \le m} D_i \models_{FDE} E_j$ for each $j \le n$. Hence $\bigwedge_{i \le m} D_i \models_{FDE} \bigwedge_{j \le n} E_j$. Similarly we can show that $\bigwedge_{j \le n} E_j \models_{FDE} \bigwedge_{i \le m} D_i$
($\Leftarrow$): Suppose $RPI(A) \equiv_{FDE} RPI(B)$. Then we have

$$A \models_{FDE} \bigwedge_{i \le m} D_i \models_{FDE} \bigwedge_{j \le n} E_j \models_{FDE} B$$

Similarly we have

$$B \models_{FDE} \bigwedge_{j \le n} E_j \models_{FDE} \bigwedge_{i \le m} D_i \models_{FDE} A$$

By the transitivity of these entailments, it follows that A and B are FDE equivalent. ∎

An immediate corollary is that standard reduction rules for CNF (DNF) conversion are RPI preserving:

Corollary 3.4.1
The following equivalences holds:

1. $\mathsf{RPI}(\neg\neg A) \equiv_{\mathsf{FDE}} \mathsf{RPI}(A)$
2. $\mathsf{RPI}(\neg(A \vee B)) \equiv_{\mathsf{FDE}} \mathsf{RPI}(\neg A \wedge \neg B)$
3. $\mathsf{RPI}(\neg(A \wedge B)) \equiv_{\mathsf{FDE}} \mathsf{RPI}(\neg A \vee \neg B)$
4. $\mathsf{RPI}(A \vee (B \wedge C)) \equiv_{\mathsf{FDE}} \mathsf{RPI}((A \vee B) \wedge (A \vee C))$

The minimality of an RPI ensures that a certain *transitivity* property of RPI holds:

Proposition 3.4.4
For any formulae A*,* B *and* C*, if* $C \in \mathsf{RPI}(B)$ *and* $B \in \mathsf{RPI}(A)$*, then* $C \in \mathsf{RPI}(A)$

Proof:
Given that $B \in \mathsf{RPI}(A)$, B must be a clause and thus $B \in \mathsf{RPI}(B)$ holds trivially. So if $C \in \mathsf{RPI}(B)$, $B \equiv_{\mathsf{FDE}} C$ follows immediately from definition (3.4.2). Hence $C \in \mathsf{RPI}(A)$.
∎

Just as Belnap's replacement rules can be used as a basis for defining the closure operators C_B and C_B^+, RPIs too can be used as a basis for defining certain Tarskian closure operators:

Definition 3.4.3
For any A *and* Γ*, define*

$$\begin{array}{rcl} C_{\mathsf{RPI}}(A) & = & \mathsf{RPI}(A) \cup \{A\} \\ C_{\mathsf{RPI}}(\Gamma) & = & \{B \in C_{\mathsf{RPI}}(A) \mid A \in \Gamma\} \\ C^+_{\mathsf{RPI}}(\Gamma) & = & \bigcup\limits_{\Delta \subseteq_{fin} C_{\mathsf{RPI}}(\Gamma)} \{B \mid B \equiv_{\mathsf{FDE}} \bigwedge \Delta\} \end{array}$$

Proposition 3.4.5
C_{RPI} *and* C^+_{RPI} *are both Tarskian closure operators. Moreover,* C^+_{RPI} *is an* E*-equivalent (*U*-equivalent) extension of* C_{RPI}*.*

Proof:
Reflexivity: trivial since $A \in C_{\mathsf{RPI}}(A)$ for every $A \in \Gamma$.

Monotonicity: Assume $\Gamma \subseteq \Delta$, then if $B \in C_{\mathsf{RPI}}(\Gamma)$, there must exist some $A \in \Gamma$ such that $B \in C_{\mathsf{RPI}}(A)$. But $A \in \Delta$ holds, so $B \in C_{\mathsf{RPI}}(\Delta)$ as required.

Idempotence: $C_{\mathsf{RPI}}(\Gamma) \subseteq C_{\mathsf{RPI}}(C_{\mathsf{RPI}}(\Gamma))$ is implied by the monotonicity of C_{RPI} above. For $C_{\mathsf{RPI}}(C_{\mathsf{RPI}}(\Gamma)) \subseteq C_{\mathsf{RPI}}(\Gamma)$, we note that proposition (3.4.4) gives us the transitivity property of RPI:

$$\begin{aligned}
D \in C_{\mathsf{RPI}}(C_{\mathsf{RPI}}(\Gamma)) &\implies \exists A \in C_{\mathsf{RPI}}(\Gamma):\ D \in C_{\mathsf{RPI}}(A)\\
&\implies \exists B \in \Gamma : A \in C_{\mathsf{RPI}}(B)\\
&\implies D \in C_{\mathsf{RPI}}(B)\\
&\implies D \in C_{\mathsf{RPI}}(\Gamma)
\end{aligned}$$

Reflexivity and monotonicity for C^+_{RPI} are straightforward. For idempotence, we verify that $C^+_{\mathsf{RPI}}(C^+_{\mathsf{RPI}}(\Gamma)) \subseteq C^+_{\mathsf{RPI}}(\Gamma)$:

$$\begin{aligned}
A \in C^+_{\mathsf{RPI}}(C^+_{\mathsf{RPI}}(\Gamma)) &\implies \exists C_1, \ldots, C_i \in C_{\mathsf{RPI}}(C^+_{\mathsf{RPI}}(\Gamma)):\\
&\qquad A \equiv_{\mathsf{FDE}} C_1 \wedge \ldots \wedge C_i\\
&\implies \exists D_1, \ldots, D_i \in C^+_{\mathsf{RPI}}(\Gamma):\\
&\qquad \forall j \leq i,\ C_j \in C_{\mathsf{RPI}}(D_j)\\
&\implies \forall j \leq i,\ \exists E^1_j, \ldots, E^m_j \in C_{\mathsf{RPI}}(\Gamma):\\
&\qquad D_j \equiv_{\mathsf{FDE}} E^1_j \wedge \ldots \wedge E^m_j\\
&\implies \forall j \leq i,\ \exists F^1_j, \ldots, F^m_j \in \Gamma:\\
&\qquad E^1_j \in C_{\mathsf{RPI}}(F^1_j), \ldots, E^m_j \in C_{\mathsf{RPI}}(F^m_j)\\
&\implies \forall j \leq i,\ C_j \in C_{\mathsf{RPI}}(E^1_j \wedge \ldots \wedge E^m_j)\\
&\implies \forall j \leq i,\ \exists k:\ C_j \equiv_{\mathsf{FDE}} E^k_j\\
&\implies \forall j \leq i,\ \exists F^k_j \in \Gamma:\ C_j \in C_{\mathsf{RPI}}(F^k_j)\\
&\implies C_1, \ldots, C_i \in C_{\mathsf{RPI}}(\Gamma)\\
&\implies A \in C^+_{\mathsf{RPI}}(\Gamma)
\end{aligned}$$

To show that C^+_{RPI} is an E-equivalent extension of C_{RPI}, we need to show that

1. C_{RPI} and C^+_{RPI} have the same E-consequences.

2. For any Γ, $C_{RPI}(\Gamma) \subseteq C^{+}_{RPI}(\Gamma)$

3. For any Γ, $C^{+}_{RPI}(C_{RPI}(\Gamma)) = C^{+}_{RPI}(\Gamma)$

(2) is trivial. For (1) we note that any FDE equivalent formula are also classically equivalent, so an argument similar to lemma (3.3.2) suffices to show that C^{+}_{RPI} is an E-equivalent (U-equivalent) extension of C_{RPI}. Finally we verify that $C^{+}_{RPI}(C_{RPI}(\Gamma)) = C^{+}_{RPI}(\Gamma)$:

($\supseteq$): Trivial since C_{RPI} and C^{+}_{RPI} are both Tarskian closure operators.

($\subseteq$): We note that C_{RPI} is idempotent.

$$\begin{aligned} A \in C^{+}_{RPI}(C_{RPI}(\Gamma)) \quad &\Longrightarrow \quad \exists B_1, \ldots, B_i \in C_{RPI}(C_{RPI}(\Gamma)) : \\ &\quad\quad A \equiv_{FDE} (B_1 \wedge \ldots \wedge B_i) \\ &\Longrightarrow \quad \exists B_1, \ldots, B_i \in C_{RPI}(\Gamma) : \\ &\quad\quad A \equiv_{FDE} (B_1 \wedge \ldots \wedge B_i) \\ &\Longrightarrow \quad A \in C^{+}_{RPI}(\Gamma) \end{aligned}$$

∎

We note that definition (3.4.3) makes use of $RPI(A)$ for each A in a given set Γ, but $\bigcup_{A \in \Gamma} RPI(A)$ need not be an independent set of RPIs. In particular redundant information can be spread across an entire set of formula. This motivates the following alternative definition:

Definition 3.4.4
For any Γ and any clause C, we define $C \in RPI^{*}(\Gamma)$ iff

1. *for some $A \in \Gamma$, $A \models_{FDE} C$ and*

2. *for any $B \in \Gamma$ and clause D, if $B \models_{FDE} D$ and $D \models_{FDE} C$, then $C \models_{FDE} D$*

For any Γ,

$$C^{*}_{RPI}(\Gamma) = RPI^{*}(\Gamma) \cup \Gamma$$

Membership for C^{*}_{RPI} is clearly more stringent then C_{RPI} – a clause C is in $RPI^{*}(\Gamma)$ only if C is entailed by some member of Γ and no other member of Γ entails a stronger clause. This definition is similar to definition (3.4.2) for the RPI's of an individual formula. However, C^{*}_{RPI} is not a closure operator in Tarski's sense. Although both reflexivity and idempotence remain intact, C^*_{RPI} does not have the usual monotonicity property.

Example 3.4.5
$\Gamma = \{p \wedge (q \vee r)\}$, $p \in C^*_{RPI}(\Gamma)$ and $q \vee r \in C^*_{RPI}(\Gamma)$. But $q \vee r \notin C^*_{RPI}(\Gamma')$ where $\Gamma' = \{p \wedge (q \vee r), q\}$

The failure of monotonicity should not be regarded as a defect of C^*_{RPI}. Arguably, *implicit information* need not always increase monotonically with respect to supersets; C^*_{RPI} is a possible candidate for specifying the content of a given set of logical expressions. To illustrate the difference between C^*_{RPI} and C_{RPI} consider the following example:

Example 3.4.6
$\Gamma = \{p, (r \wedge \neg r) \wedge (p \vee q), \neg p\}$

Since $p \in C^*_{RPI}(\Gamma)$ we have $p \vee q \notin C^*_{RPI}(\Gamma)$. However $p \vee q \in C_{RPI}(\Gamma)$ given that $p \vee q \in RPI((r \wedge \neg r) \wedge (p \vee q))$. Note that in example (3.4.6) q is an E-consequence of $C_{RPI}(\Gamma)$ but not an E-consequence of $C^*_{RPI}(\Gamma)$. In general, C^*_{RPI} does not yield the same E-consequence (U-consequence) as C_{RPI}.

Proposition 3.4.6
For any Γ*,*

1. $C^*_{RPI}(\Gamma) \subseteq C_{RPI}(\Gamma)$

2. $C_{PRI}(C^*_{RPI}(\Gamma)) = C_{RPI}(\Gamma)$

3. $C^*_{RPI}(C_{RPI}(\Gamma)) = C_{RPI}(\Gamma)$

Proof:
For (1) it suffices to observe that $RPI^*(\Gamma) \subseteq \bigcup_{A \in \Gamma} RPI(A)$.

(2$\supseteq$): Since $\Gamma \subseteq C^*_{RPI}(\Gamma)$, we have $C_{RPI}(\Gamma) \subseteq C_{RPI}(C^*_{RPI}(\Gamma))$ by the monotonicity of C_{RPI}.

(2$\subseteq$): From (1) we have $C^*_{RPI}(\Gamma) \subseteq C_{RPI}(\Gamma)$ so by the monotonicity of C_{RPI} it follows that $C_{RPI}(C^*_{RPI}(\Gamma)) \subseteq C_{RPI}(C_{RPI}(\Gamma))$. By the idempotence of C_{RPI} we have $C_{RPI}(C^*_{RPI}(\Gamma)) \subseteq C_{RPI}((\Gamma)$

(3): Trivial from (1). ∎

Returning to examples (3.4.1) and (3.4.2), Belnap's replacement rules are complete with respect to the given A and B in these examples, i.e. $C_{RPI}(A) \subset C_B(A)$ and $C_{RPI}(B) \subset C_B(B)$, but the generated implicates are not all prime. So Belnap's replacement rules are unsound with respect to relevant prime implicates. In the general case, Belnap's replacement rules are not complete since they are insufficient to transform formulae into clausal form. Clearly for clause reduction we need the additional rule, $\vdash \neg(B \wedge C) \leftrightarrow (\neg B \vee \neg C)$, to distribute negation over conjunction. However C^+_B is complete with respect to RPI's, i.e. for any Γ, we have $C_{RPI}(\Gamma) \subseteq C^+_B(\Gamma)$. We summarise the relationships of these closure operators in figure (3.8).

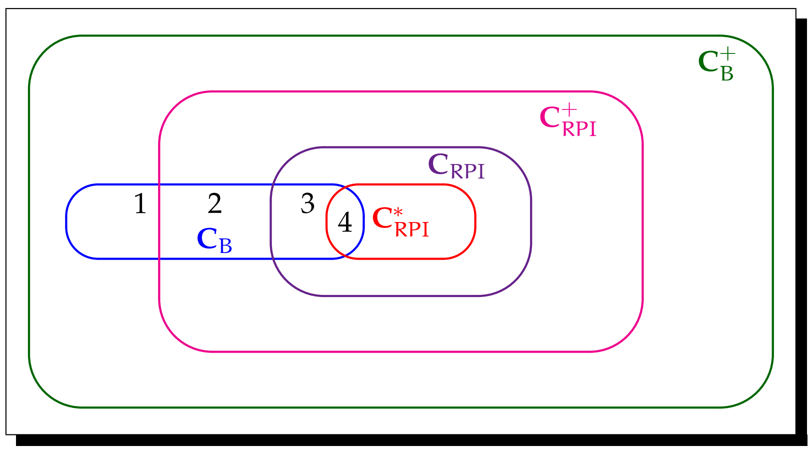

Figure 3.8: Relationships between Closure Operators

To illustrate consider $\Gamma = \{(p \wedge q) \vee (p \wedge r), \neg(p \wedge q) \wedge s\}$. Clearly, $p \vee r \in C_B(\Gamma)$ but $p \vee r$ is not an RPI, so $p \vee r \notin C^+_{RPI}(\Gamma)$. Region (1) is non-empty. Moreover $\neg(p \wedge q) \in C_B(\Gamma)$ but $\neg p \vee \neg q \in C^+_{RPI}(\Gamma)$, so region (2) is non-empty. Example (3.4.6) shows that region (3) is non-empty and with minor modification it can show that region (4) is also non-empty. To see that $C^+_{RPI} \subseteq C^+_B$, it suffices to note that

Proposition 3.4.7
For any clause D *and formula* A*, if* $D \in RPI(A)$*, then* $E \equiv_{FDE} D$ *for some* $E \in C_B(A)$.

Proof:
We note that using arguments similar to the proofs of lemma (3.3.1) and proposition (3.3.1), we can show that any A is FDE-equivalent to the conjunction of the leaves of the Belnap's tree $\mathcal{T}^A$, i.e. $le(\mathcal{T}^A) \equiv_{FDE} RPI(A)$. Moreover, each leaf of a Belnap's tree is conjunction free in the sense that each leaf is FDE-equivalent to a clause. Hence if $D \in RPI(A)$, there must be a leaf E of $\mathcal{T}^A$ such that $D \equiv_{FDE} E$. ∎

We should point out that in adopting the use of either C_{RPI} or C^{+}_{RPI} for capturing the informational content of a formula, there is no guarantee that conjunction elimination is a sound strategy for generating RPIs. In general $RPI(A) \cup RPI(B) \neq RPI(A \wedge B)$.

Example 3.4.7
$A = p \wedge (p \vee q)$

In example (3.4.7), it is clear that $RPI(p \wedge (p \vee q)) \subset RPI(p) \cup RPI(p \vee q)$. The containment here is proper. However, we do have the containment $RPI(A \wedge B) \subseteq RPI(A) \cup RPI(B)$ in the general case.

Lemma 3.4.1
For any clause C *and any formula* A *and* B*, if* $A \models_{FDE} C$*, then* $A \wedge B \models_{FDE} C$.

Proof:
Again we use the ambi-valuation of Dunn ([66]) to prove our claim. Assume that $A \models_{FDE} C$. Then we have the implication $1 \in v(A) \Rightarrow 1 \in v(C)$ for any standard 4-valued valuation v of FDE. Consider an arbitrary v' such that $1 \in v'(A \wedge B)$. Then it follows that $1 \in v'(A)$ and $1 \in v'(B)$. So on v' in particular, $1 \in v'(C)$. Since v' was arbitrary, we have $A \wedge B \models_{FDE} D$ as required. ∎

Proposition 3.4.8
For any A *and* B, $RPI(A \wedge B) \subseteq RPI(A) \cup RPI(B)$.

Proof:
Assume that for an arbitrary clause D we have $D \in RPI(A \wedge B)$ but $D \notin RPI(A) \cup RPI(B)$. Then we have $A \wedge B \models_{FDE} D$ but $D \notin RPI(A)$ and $D \notin RPI(B)$. Then there are 4 cases to consider:

(case 1) $A \not\models_{FDE} D$ and $B \models_{FDE} D$ but D is not prime for B: it follows that there exists a $D_0 \in RPI(B)$ such that $B \models_{FDE} D_0$ and $D_0 \models_{FDE} D$ but $D \not\models_{FDE} D_0$. By lemma (3.4.1), we have $A \wedge B \models_{FDE} D_0$. But given that $D \in RPI(A \wedge B)$ and $D_0 \models_{FDE} D$, $D \models_{FDE} D_0$ holds. This is a contradiction.

(case 2) $B \not\models_{FDE} D$ and $A \models_{FDE} D$ but D is not prime for A: the proof is similar to case 1 with B replaced with A throughout.

(case 3) Both $A \models_{FDE} D$ and $B \models_{FDE} D$, but D is prime for neither A nor B: the argument in case (1) suffices to show that case 3 is impossible.

(case 4) $A \not\models_{FDE} D$ and $B \not\models_{FDE} D$: we make use of the equivalence between FDE and *tautological entailment* as described in Anderson and Belnap ([7]). Since $A \not\models_{FDE} D$ and $B \not\models_{FDE} D$, for any arbitrary but fixed $DNF = C_1^A \vee \ldots \vee C_m^A$ of A and $DNF = C_1^B \vee \ldots \vee C_n^B$ of B, there exist some $i \leq m$, and some $j \leq n$ such that $C_i^A \not\models_{FDE} D$ and $C_j^B \not\models_{FDE} D$. Denote the set of literals occurring in C_i^A as $lit(C_i^A)$. We have $lit(C_i^A) \cap lit(D) = \emptyset$ and $lit(C_j^B) \cap lit(D) = \emptyset$. Hence $(lit(C_i^A) \cup lit(C_j^B)) \cap lit(D) = \emptyset$. Now consider the formula

$$E = \bigvee_{1 \leq k \leq m,\ 1 \leq l \leq n} (C_k^A \wedge C_l^B)$$

Clearly, E is a DNF of $A \wedge B$. Since $lit(C_i^A \wedge C_j^B) = (lit(C_i^A) \cup lit(C_j^B))$, we note that $lit(C_i^A \wedge C_j^B) \cap lit(D) = \emptyset$. We define a 4-valued assignment v on the set of propositional atoms as follows:

$$\begin{cases} 0 \in v(p) \text{ and } 1 \notin v(p) & \text{if } \neg p \in lit(C_i^A \wedge C_j^B) \text{ and} \\ & p \notin lit(C_i^A \wedge C_j^B) \\ 1 \in v(p) \text{ and } 0 \notin v(p) & \text{if } p \in lit(C_i^A \wedge C_j^B) \text{ and} \\ & \neg p \notin lit(C_i^A \wedge C_j^B) \\ 1 \in v(p) \text{ and } 0 \in v(p) & \text{if } p \in lit(C_i^A \wedge C_j^B) \text{ and} \\ & \neg p \in lit(C_i^A \wedge C_j^B) \\ 1 \notin v(p) \text{ and } 0 \notin v(p) & \text{otherwise} \end{cases}$$

Clearly $1 \in v(C_i^A \wedge C_j^B)$ and hence $1 \in v(E)$ but by the disjointness of $lit(C_i^A \wedge C_j^B)$ and $lit(D)$, $1 \notin v(D)$. Hence $A \wedge B \not\models_{FDE} D$. But this contradicts the initial assumption that $D \in RPI(A \wedge B)$. ∎

3.4.2 Algorithmic Considerations

Proposition (3.4.8) shows that in terms of using replacement rules in the style of [$*$] or [#] for eliminating conjunctions, the RPIs of a child node need not be the RPIs of the root node. So although corollary (3.4.1) shows that the standard reduction method for CNF conversion is indeed complete for generating RPIs, there is no guarantee that the clauses obtained are indeed independent. Checking for clause subsumption seems unavoidable and indeed critical when redundant information is presented. However when combined with a clause subsumption check, the standard CNF conversion algorithm can provide a sound and complete algorithm for generating RPIs.

Algorithm 3.4.1 RPI Generation

Require: input $A \in \Phi$
Ensure: output $S = RPI(A)$
1: convert A into $CNF(A)$ using the standard reduction method
2: for each $C \in CNF(A)$, $S := S \cup \{C\}$ if C is relevant prime, else $S := S$.
3: return S

Algorithm (3.4.1) is a naive method for generating RPIs. It first generates a set of relevant implicates of A and then prunes the set by removing all non-prime implicates. Clearly we have $CNF(A) \equiv_{FDE} RPI(A)$ given corollary (3.4.1). So completeness is ensured in step (1) provided that step (2) does not remove implicates that are also prime (and clearly it doesn't). Although the clause subsumption check may be deployed earlier while $CNF(A)$ is generated, in the worst case the size of $CNF(A)$ can be exponentially related to the size of A, e.g. if $A = (p_1 \wedge p_2) \vee \ldots \vee (p_{2n-1} \wedge p_{2n})$, there are 2^n clauses in the corresponding CNF. Our problem is inherently difficult computationally.

3.4.2.1 PRI via Classical PI Generation

In what follows, we'll present an alternative algorithm for generating $RPI(A)$ based on ideas from Ramesh *et al* [147; 145; 146] and Arieli and Denecker [13; 14]. The main idea here is to avoid the expensive CNF conversion by using *negated normal form* (NNF) instead. Once a formula A is converted into $NNF(A)$, we'll make use of Arieli and Denecker's *splitting transform* to convert $NNF(A)$ into a positive (i.e. negation free) formula $\widehat{NNF(A)}$.[1] The conversion will preserve our problem in the sense that for any clause D, $D \in RPI(A)$ iff $\widehat{D} \in PI(\widehat{NNF(A)})$. So in effect our problem is transformed into the classical problem of prime implicate generation for a positive NNF formula. The algorithm of [146; 145] can thus be invoked to generate the required PI's via the use of the corresponding *semantic graph*. Before we present the algorithm, we need to present some of the main definitions.

[1] We note that Besnard and Schaub [38] employed the same transform for defining signed systems of paraconsistent reasoning.

Definition 3.4.5

1. *A formula* A *is in negated normal form (*NNF*) iff no complex subformula of* A *is in the scope of a negation, i.e. only atomic formulae are within the scope of a negation operator.*

2. *Let* $NNF(A)$ *denotes the negated normal form of* A*. Then the splitting transform of* $NNF(A)$*, denoted by* $\widehat{NNF(A)}$*, is the formula obtained by uniformly substituting every unnegated atom* p_i *occurring in* $NNF(A)$ *with a new (signed) atom* p_i^+ *and every negated atom* $\neg p_i$ *in* $NNF(A)$ *with a new (signed) atom* p_i^- *[13; 14]. If* $B = \widehat{A}$ *for some* A*, then we define the inverse of splitting transform* $\overline{B}$ *as the formula obtained by uniformly substituting every signed atom* p_i^+ *with literal* p_i *and every signed atom* p_i^- *with literal* $\neg p_i$*, i.e.* $\overline{\widehat{A}} = A$

3. *Let* v *be an arbitrary* 4*-valued assignment and* $NNF(A)$ *an arbitrary* NNF *formula. Then* $\widehat{v}$ *is the* 2*-valued (classical) assignment defined as follows:*

   - *For all* p_i^+ *and* p_i^- *occurring in* $\widehat{NNF(A)}$*,* $\widehat{v}(p_i^+) = 1$ *iff* $1 \in v(p_i)$ *and* $\widehat{v}(p_i^-) = 1$ *iff* $0 \in v(p_i)$*.*

We note that both the splitting transform and $\widehat{v}$ are well defined and do not depend on A. The following are consequences of definition (3.4.5):

Proposition 3.4.9

1. *Let* v *be an arbitrary* 4*-valued assignment and* $NNF(A)$ *be an arbitrary* NNF *formula. Let* $\widehat{v}$ *be a* 2*-valued assignment as defined in (3) of definition (3.4.5). Then* $1 \in v(NNF(A))$ *iff* $\widehat{v}(\widehat{NNF(A)}) = 1$ *(cf. Lemma (3.1) of [13])*

2. *For any* A *and* B*,* $A \models_{FDE} B$ *iff* $\widehat{NNF(A)} \models \widehat{NNF(B)}$ *(cf. Theorem (3.1) of [13]).*

3. *For any clause* D*,* $D \in RPI(A)$ *iff* $\widehat{D} \in PI(\widehat{NNF(A)})$*.*

4. *The problem of relevant prime implicate generation is polynomially reducible to classical prime implicate generation.*

Proof:

We note that (2) is a simple corollary of (1). For (1), we use an induction on the structure of $NNF(A)$. There are two base cases with either $NNF(A) = p_i$ or $NNF(A) = \neg p_i$. In the former case we have $1 \in v(p_i) \Leftrightarrow \widehat{v}(p_i^+) = 1$ given by the definition of $\widehat{v}$. In the later case we have $1 \in v(\neg p_i) \Leftrightarrow 0 \in v(p_i) \Leftrightarrow \widehat{v}(p_i^-) = 1$.

For the induction case we have either $\mathsf{NNF}(A) = B \wedge C$ or $\mathsf{NNF}(A) = B \vee C$. We note that both B and C must be in NNF form and hence the induction hypothesis applies. Thus we have $1 \in \nu(B) \Leftrightarrow \widehat{\nu}(\widehat{B}) = 1$ and $1 \in \nu(C) \Leftrightarrow \widehat{\nu}(\widehat{C}) = 1$. So $1 \in \nu(B \wedge C) \Leftrightarrow [1 \in \nu(B) \text{ and} 1 \in \nu(C)] \Leftrightarrow [\widehat{\nu}(\widehat{B}) = 1 \text{ and } \widehat{\nu}(\widehat{C}) = 1] \Leftrightarrow \widehat{\nu}(\widehat{B \wedge C}) = 1$. The case for $B \vee C$ is similar.

(3 $\Rightarrow$): Since $\mathsf{NNF}(A) \equiv_{\mathsf{FDE}} A$ it suffices for us to consider an arbitrary $D \in \mathsf{RPI}(\mathsf{NNF}(A))$. Then by (2) above we have $\widehat{\mathsf{NNF}(A)} \models \widehat{D}$. This shows that $\widehat{D}$ is an implicate of $\widehat{\mathsf{NNF}(A)}$. Toward a contradiction, suppose $\widehat{D}$ is not prime. Then there exists a clause C such that $\widehat{\mathsf{NNF}(A)} \models C$ and $C \models \widehat{D}$ but $\widehat{D} \not\models C$. But $\widehat{\mathsf{NNF}(A)}$ is negation free and thus neither C nor $\widehat{D}$ are the empty clause, nor are they tautologies. Hence there must be a C' such that $\mathsf{NNF}(A) \models_{\mathsf{FDE}} C'$ where $\widehat{C'} = C$. But then we have $C' \models_{\mathsf{FDE}} D$ but $D \not\models_{\mathsf{FDE}} C'$. This contradicts the primeness of D. Hence $\widehat{D} \in \mathsf{PI}(\widehat{\mathsf{NNF}(A)})$ as required.

(3 $\Leftarrow$): Suppose that $D \notin \mathsf{RPI}(\mathsf{NNF}(A))$. Then either $\mathsf{NNF}(A) \not\models_{\mathsf{FDE}} D$ or D is not prime. In the former case, $\widehat{\mathsf{NNF}(A)} \not\models \widehat{D}$ follows immediately from (2). So suppose D is relevant implicate of $\mathsf{NNF}(A)$ but is not prime. Then there exists a C such that $\mathsf{NNF}(A) \models_{\mathsf{FDE}} C$ and $C \models_{\mathsf{FDE}} D$ but $D \not\models_{\mathsf{FDE}} C$. By (3$\Rightarrow$) and (2) above it follows that $\widehat{C}$ is prime implicate of $\widehat{\mathsf{NNF}(A)}$ but $\widehat{D} \not\models \widehat{C}$. Hence $\widehat{D} \notin \mathsf{PI}(\widehat{\mathsf{NNF}(A)})$ as required.

(4): We note that both NNF conversion and the splitting transform are linearly related to the input formula. Hence by (3) above, the claim follows. ∎

3.4.2.2 Semantic Graphs

In [147], a non-clausal approach is developed to compute classical prime implicates (implicants). Their basic motivation is to avoid the computational overhead of CNF/DNF based approaches to prime implicate (implicant) generation that are known to be exponential. Their general strategy is to use an alternative graph theoretic representation of NNF formulae and to reduce the search for prime implicates (implicants) to a search of un-subsumed d-paths (c-paths) in the graph. Although the associated decision problems are known to be **NP**-hard and **NP**-complete (theorem 8 and theorem 9 of [146]), their theoretical framework is both sound and elegant. Their experimental results also show that for a certain class of formulae that are hard for CNF − DNF based methods, their algorithms show significant improvement. We present some of their central definitions and results here.

Definition 3.4.6
Given an arbitrary formula A *in* NNF*, the semantic graph associated with* A*,* $\mathcal{G}_A = (\mathcal{N}, \mathcal{C}, \mathcal{D})$*, is a triple where the nodes* $\mathcal{N}$ *are the literals of* $\mathrm{NNF}(A)$*,* $\mathcal{C}$ *is the set of* c*-arcs and* $\mathcal{D}$ *is the set of* d*-arcs. A* c*-arc (*d*-arc) is a conjunction (disjunction) of two semantic subgraphs of* $\mathcal{G}_A$*. The notion of a subgraph is defined recursively in the standard way where all elements of* $\mathcal{N}$ *are subgraphs of* $\mathcal{G}_A$ *and all complex subgraphs are built up using* c*-arcs and* d*-arcs.*

Essentially a semantic graph is an alternative two dimensional representation of an NNF formula. Following the standard convention, we display c-arcs vertically and d-arcs horizontally and continue to overload our symbols '$\wedge$' for c-arc and '$\vee$' for d-arc. Both c-arcs and d-arcs are obviously associative and commutative. We'll use the convention that a d-arc is displayed horizontally in the same order as the corresponding disjunction from left to right, whereas a c-arc is displayed vertically with the right conjunct *below* the left conjunct (see figure (3.9)). To increase readability we may display a semantic graph with boxes around certain subgraphs. Like many other graph based representations, the notion of a *path* plays a central role in graph based computation.

Example 3.4.8
$A = [(\neg p \wedge q) \vee (r \vee s)] \wedge [\neg q \vee (t \wedge p)]$

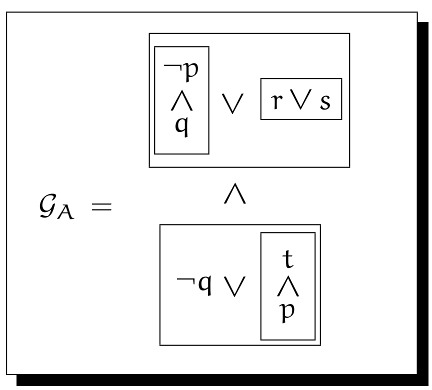

Figure 3.9: Semantic Graph of A

Definition 3.4.7
Let $\mathcal{G}_A$ *be a semantic graph. Let* X *and* Y *be subgraphs of* $\mathcal{G}_A$ *and let* a *and* b *be nodes in* X *and* Y *respectively. Then* a *and* b *are said to be* α*-connected iff* (X, Y) *is an* α*-arc* $(\alpha = c, d)$*. An* α*-link is a complementary (e.g.* p *and* $\neg p$*) pair of* α*-connected nodes.*

A partial α-path through $\mathcal{G}_A$ is a multi-set of pairwise α-connected nodes of $\mathcal{G}_A$. An α-path through $\mathcal{G}_A$ is a partial α-path that is maximal, i.e. it has no proper extension that is also a partial α-path through $\mathcal{G}_A$.

In example (3.4.8), $\{\neg p, q, t, p\}$, $\{\neg p, q, \neg q\}$, $\{r, \neg q\}$, $\{s, \neg q\}$, $\{r, t, p\}$, and $\{s, t, p\}$ are all c-paths. On the other hand, $\{\neg p, r, s\}$, $\{q, r, s\}$, $\{\neg q, t\}$ and $\{\neg q, p\}$ are all d-paths. Although our definition of an α-path is strictly stated in terms of *multisets*, it would be convenient to continue to use set theoretic representation for paths. We say that one α-path is subsumed by another α-path with the understanding that the set of nodes of one path is contained in the set of nodes of the other. We write $\mathtt{lit}(\mathcal{P})$ to denote the set of literals (nodes) of the path $\mathcal{P}$. Intuitively, a linkless c-path through $\mathcal{G}_A$ corresponds to a model of A, whereas the set formed by complementing all literals occurring in a linkless d-path through $\mathcal{G}_A$ corresponds to a counter-model of A, i.e. a truth assignment which falsifies A.

Some of the most important equivalence preserving operations on semantic graph are *path dissolution* operations. Strictly speaking there are two types of path dissolution operations corresponding to c-path and d-path dissolution. Full details and definitions are in [146], but roughly the idea of path dissolution is to select a link and then restructure a semantic graph so that any α-path with the link is removed. A special case of c-path dissolution is just the standard resolution rule. A semantic graph is said to be a *full dissolvent* if all of its α-paths are linkless. Dissolution for c-paths is strongly complete in the sense that any sequence of dissolution steps will terminate in a c-linkless semantic graph. In the event that a semantic graph corresponds to an unsatisfiable formula, repeat application of c-path dissolution will terminate with the empty graph $(\emptyset, \emptyset, \emptyset)$. Similarly dissolution for d-paths will also terminate in a d-linkless semantic graph. And in the event that the semantic graph corresponds to a tautology, repeat application of d-path dissolution will also terminate with the empty semantic graph. In this context the empty graph is ambiguous – it represents both $\top$ and $\bot$ in classical logic. In figure (3.9) for instance, $\mathcal{G}_A$ is d-linkless but not c-linkless. The interest in linkless graphs lies in the following theorem:

Theorem 3.4.1

([146] theorem 3) In any nonempty semantic graph in which no c*-path (*d*-path) contains a link, every implicate (implicant) of the corresponding formula is subsumed by some* d*-path (*c*-path).*

Given the fact that any $\widehat{\mathsf{NNF}(\mathsf{A})}$ is negation free, any semantic graph $\mathcal{G}_{\widehat{\mathsf{NNF}(\mathsf{A})}}$ must be linkless (i.e. neither c-paths nor d-paths contain any links). Hence by (3) of proposition (3.4.9) and theorem (3.4.1) above, we have the following immediate corollary:

Corollary 3.4.2
Let A *be any formula and* D *be a clause, then* $\mathsf{D} \in \mathsf{RPI}(\mathsf{A})$ *iff* $\mathtt{lit}(\widehat{\mathsf{D}})$ *is subsumed by some* d*-path through the semantic graph* $\mathcal{G}_{\widehat{\mathsf{NNF}(\mathsf{A})}}$

Corollary (3.4.2) provides us with the basis to incorporate the **PI** algorithm described in [145] to generate RPIs. Moreover given (4) of proposition (3.4.9), the computational overhead of the splitting transform will not adversely affect the **PI** algorithm. However corollary (3.4.2) only warrants completeness but not soundness of the method. In order to capture the exact RPI's of a given A we need to define the largest subset of d-paths of $\mathcal{G}_\mathsf{A}$ that is *minimal*:

Definition 3.4.8
Let $\mathcal{G}_\mathsf{A}$ *be a non-empty semantic graph without* c*-links, then* $\pi(\mathcal{G}_\mathsf{A})$ *is defined as follows:* $\mathcal{P} \in \pi(\mathcal{G}_\mathsf{A})$ *iff*

1. $\mathcal{P}$ *is a* d*-path through* $\mathcal{G}_\mathsf{A}$*,*
2. $\mathcal{P}$ *is linkless (i.e.* $\mathcal{P} \neq \top$*) and*
3. *For all* d*-paths* $\mathcal{Q}$ *through* $\mathcal{G}_\mathsf{A}$*,* $\mathtt{lit}(\mathcal{Q}) \not\subset \mathtt{lit}(\mathcal{P})$*, i.e.* $\mathcal{P}$ *is not subsumed by any other* d*-path through* $\mathcal{G}_\mathsf{A}$*.*

In definition (3.4.8), $\pi(\mathcal{G}_\mathsf{A})$ captures exactly the set of prime implicates of A (theorem 6 [145]). Thus for any A we have

$$\mathsf{RPI}(\mathsf{A}) = \{\overline{\bigvee \mathtt{lit}(\mathcal{P})} : \mathcal{P} \in \pi(\mathcal{G}_{\widehat{\mathsf{NNF}(\mathsf{A})}})\}$$

where $\bigvee \mathtt{lit}(\mathcal{P})$ is a disjunction formed with the literals of the path $\mathcal{P}$ and $\overline{\bigvee \mathtt{lit}(\mathcal{P})}$ is the inverse of the splitting transform of $\bigvee \mathtt{lit}(\mathcal{P})$. In short we have the following alternative algorithm for generating the RPIs for any given A:

In the original **PI** algorithm for computing classical prime implicates (implicants) of an arbitrary NNF formula, additional steps are required to first dissolve a semantic graph into a full dissolvent (with respect to c-paths). But this is clearly not required for us since any $\mathcal{G}_{\widehat{\mathsf{NNF}(\mathsf{A})}}$ is linkless. Similarly **PI** is not required to check whether a d-path contains links ((2) of definition (3.4.8)). The only real work to be done is to check for

Algorithm 3.4.2 RPI Generation via splitting transform and semantic graph

Require: input $A \in \Phi$
Ensure: output $S = \mathrm{RPI}(A)$

```
convert A into NNF(A) using the standard reduction method
convert NNF(A) to \widehat{NNF(A)}
call PI to compute π(G_{\widehat{NNF(A)}})
S := {\overline{\bigvee lit(P)} : P ∈ π(G_{\widehat{NNF(A)}})}
return S
```

path subsumption ((3) of definition (3.4.8)). But unfortunately the path subsumption check is the most computationally intensive task of the algorithm.

The basic idea of **PI** is to recursively traverse $\mathcal{G}_A$ from a left-to-right, bottom-up manner while partial paths are computed along the way. Both tautologies and subsumed paths are eliminated as they are encountered. Each recursive call to **PI** takes a collection of (possibly empty) sets of d-paths $\{\mathcal{P}_i | \, i \in I\}$ and a full dissolvent semantic (sub)graph $\mathcal{G}_B$ as inputs and correctly computes the maximal set of minimal d-paths, $\pi(\mathcal{G}_B)$ (theorem 8 [145]). The first call to **PI** will take $(\{\emptyset\}, \mathcal{G}_A)$ as input and then subsequent recursive calls will move **PI** progressively towards the left bottom-most node of $\mathcal{G}_A$. If a subgraph $\mathcal{G}_B$ is of the form $(X, Y)_c$ then **PI** will first attempt to find solutions for Y and then solutions for X while making sure that any subsumed d-path in X is removed. If $\mathcal{G}_B$ is of the form $(X, Y)_d$, then **PI** will attempt to find solution for X first and then extend all d-paths through X into Y. The set of d-paths $\{\mathcal{P} : \; i \in I\}$ at each recursive call is the set of non-tautological and unsubsumed d-paths that have been traversed by **PI**. Initially $\{\mathcal{P} : \; i \in I\}$ is empty until **PI** meets the left bottom-most node of $\mathcal{G}_A$. Whenever **PI** is called for a node, it will attempt to extend all current d-paths with the new node unless the complement of the node is already in a path (thereby avoiding tautologous paths). This is done repeatedly until all nodes in the graph are visited by **PI** at least once.

To illustrate consider $\mathcal{G}_A$ in figure (3.10). After the initial recursive calls, **PI** will meet the left bottom most node of $\mathcal{G}_A$. In this case it is n_1. **PI** will update its current set of d-paths to $\{\{n_1\}\}$ then go right to n_2 and see if the single d-path $\{n_1\}$ can be extended. Then **PI** will continue upward to see if n_3 can be added to $\{n_1\}$. At the end of each move upward along a c-path, **PI** must check for path subsumption in the current set of d-paths and remove any subsumed path. In this case if n_2 and n_3 are neither complementary literals nor identical literals, then the current set of d-paths would be $\{\{n_1, n_2\}, \{n_1, n_3\}\}$. After traversing through n_1, n_2 and n_3, **PI** will again

Algorithm 3.4.3 PI for computing $\pi(\mathcal{G}_A)$

Require: input $(\mathrm{paths}, \mathcal{G}_A)$
Ensure: output $\mathrm{paths}'' = \pi(\mathcal{G}_A)$

PI$(\mathrm{paths}, \mathcal{G}_A)$
if $\mathrm{paths} = \emptyset$ **then**
    return $\emptyset$
end if
if $\mathcal{G}_A$ is a literal **then**
    $\mathrm{paths}' := \emptyset$;
    $\mathrm{paths}'' := \emptyset$
    for all $\mathcal{P} \in \mathrm{paths}$ **do**
        if $\mathcal{G}_A \in \mathcal{P}$ **then**
            $\mathrm{paths}' = \mathrm{paths}' \cup \mathcal{P}$
        else if $\neg\mathcal{G}_A \notin \mathcal{P}$ **then**
            $\mathrm{paths}'' := \mathrm{paths}'' \cup \mathcal{P} \cup \mathcal{G}_A$
        end if
    end for
    $\mathrm{paths}'' := \mathrm{paths}' \cup (\mathrm{paths}' \setminus \{\mathcal{P} \in \mathrm{paths}' : \exists \mathcal{Q} \in \mathrm{paths}' \wedge \mathcal{Q} \subset \mathcal{P}\})$
    return paths''
else if $\mathcal{G}_A = (X, Y)_c$ **then**
    $\mathrm{paths}' :=$ **PI**(paths, Y)
    $\mathrm{paths}'' :=$ **PI**$((\mathrm{paths} \setminus \mathrm{paths}'), X)$
    $\mathrm{paths}'' := (\mathrm{paths}' \cup \mathrm{paths}'')$
            $\setminus \{\mathcal{P} \in \mathrm{paths}' | \exists \mathcal{Q} \in \mathrm{paths}'' \wedge \mathcal{Q} \subset \mathcal{P}\}$
            $\setminus \{\mathcal{P} \in \mathrm{paths}'' | \exists \mathcal{Q} \in \mathrm{paths}' \wedge \mathcal{Q} \subset \mathcal{P}\}$
    return paths''
else if $\mathcal{G}_A = (X, Y)_d$ **then**
    $\mathrm{paths}' :=$ **PI**(paths, X)
    $\mathrm{paths}'' :=$ **PI**(paths', Y)
    return paths''
end if

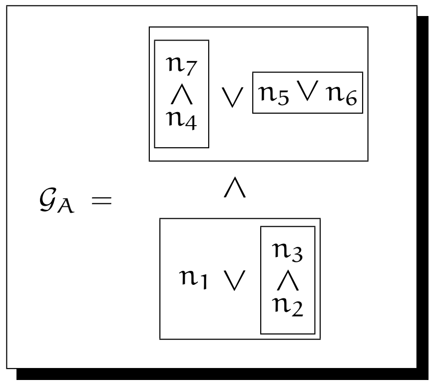

Figure 3.10: Construction of $\pi(\mathcal{G}_A)$ using **PI**

move to the left bottom most node of the upper subgraph, meeting n_4. **PI** will then attempt to extend $\{n_4\}$ with n_5 and continuing with n_6. After meeting n_6, **PI** will once again check for subsumption, if there is no subsumption then once again the current set of d-paths is updated to $\{\{n_1, n_2\}, \{n_1, n_3\}, \{n_4, n_5, n_6\}\}$. **PI** then continues with a new d-path starting with n_7 moving right toward n_5 and then n_6. Finally **PI** will complete the last subsumption check before updating the current set of d-paths one last time.

As we pointed out earlier, path subsumption checking is computationally expensive. Thus in [146], various optimisation techniques based on *anti-link operations* are introduced to restructure a semantic graph before **PI** is invoked. The restructuring involves early removal of subsumed d-paths. In some cases the improvement is dramatic – without anti-link operations **PI** will take exponential time to find the solution whereas with anti-link operations **PI** will only take polynomial time. Anti-links are essentially multiple occurrences of the same literal. If $(X, Y)_\alpha$ is an α-arc in a semantic graph and A_X and A_Y are nodes of the same literal A in X and Y respectively, then $\{A_X, A_Y\}$ is said to be an (α-) anti-link. The presence of an anti-link in a semantic graph is a necessary (but not sufficient) condition for the existence of a non-tautologous subsumed d-paths in a semantic graph (theorem 15 [146]), i.e. the occurrence of a non-tautologous subsumed d-paths implies the existence of either a c-anti-link or a d-anti-link in a semantic graph. The amount of subsumption checking required can be reduced by removing anti-links preemptively. In some cases, anti-links can be completely eliminated in a semantic graph, thereby eliminating the need for any subsumption checking. But in the general case, complete elimination of anti-links in a semantic graph is not always possible.

3.5 Conclusion

In this chapter we have seen that reasoning with inconsistent information can be divided into two distinct stages. In the first stage inconsistent information encoded in a full language can be rewritten in such a way as to facilitate the isolation of the inconsistent part of the information. In the second stage various deduction strategies based on either classical or nonclassical logics can then be applied to the rewrite. We note that Belnap's strategy of dividing reasoning into a *preprocessing* stage and a *deduction* stage is akin to a recent approach to *knowledge compilation*. In knowledge compilation a knowledge base encoded in a logical language is first complied into a target language. The complied knowledge base is then deployed during run-time query answering. The main objective of the compilation is to make on-line reasoning easier. The hope is that the time required for compilation will have an eventual payoff during run-time query answering.

With respect to preprocessing inconsistent information however, we find Belnap's suggestion of using *conjunctive containment* wanting. In particular, inconsistent information tends to interact badly with disjunctive and redundant information. Although conjunctive containment generally reduces disjunctive consequences, it is however insufficient. Our remedy is to use a relevant notion of prime implicates as the basis to both preserve information and minimise the potentially harmful disjunctive content of inconsistent information.

Chapter 4

Uncertainties and Inconsistencies

4.1 Introduction

In previous chapters, we have introduced a preservation theoretic analysis of inference. We have done so both in terms of a measurement based on coverings of a set as well as an information theoretic measurement. We have also argued that, at least with respect to the method of reasoning from maximal consistent subsets, we generally need to pay attention to the syntactic form of the premises since these measurements are essentially syntax sensitive notions. In this chapter, we approach the issue of preservation from a slightly different angle; we'll look at the problem of *uncertainties* that are transmitted from an inconsistent set of premises to the conclusion in a given inference. The problem can be stated as follows: suppose we are given an inference with an inconsistent set of premises $\{A_1, \ldots, A_m\}$ together with the conclusion B derivable in some logic L. Suppose further that we have a particular method for assigning uncertainties to each premise and to the conclusion in terms of probabilities. Furthermore let's suppose that for each $i \leq m$, the uncertainty of each premise is $\pi_i \in [0, 1]$, i.e. $U(A_i) = \pi_i$. What then is the maximum value of the uncertainty of the conclusion?

Our question is important for several reasons. Firstly, even if we have a classically valid inference with consistent premises it is not always the case that we can be completely certain about each one of the premises. The validity of an inference only guarantees that if there is no uncertainty in the premises, i.e. they are all true, then there can be no uncertainty in the conclusion. In an extreme case we may have a valid inference with a billion consistent premises and the conclusion is the conjunction of all the premises. The uncertainty of each premise may be less then one in a billion but the cumulative uncertainty of the conclusion may turn out to be prohibitively high. To determine the uncertainty maxima of the conclusion of an inference is one way to provide probabilistic assurance in uncertain reasoning – a kind of assurance that goes beyond mere deductive validity. Secondly, in the case of inferences with inconsistent premises we have a general problem of assessing the general *quality* of the conclusion. This problem cannot be solved by simply switching to a nonclassical logic. Indeed with a few notable exceptions, most proponents of paraconsistent logics have not really addressed the issue of quality control at all. Even with a seemingly innocent inference like the conjoining of premises, the resulting conclusion may turn out to be unacceptable due to high uncertainty. Finally, our problem is important in light of the remarks of Adams and Levine in [5] that

> inconsistent premise sets introduce some surprising uncertainty phenomena whose interpretations involve problems (p.432)

In this respect, the current chapter can be seen as a continuation of the unfinished work of Adams and Levine in [5]. But the theoretical foundation can be traced back to the work of Boole a century ago in [42] and has been revived, a century later, by Hailperin in [81] and more recently by Nilsson in [134] and Knight in [109; 110; 111]. We do not proclaim that the theoretical work here is particularly new; but we do hope to show that the general framework of probabilistic analysis of inferences fits well with the preservational approach to paraconsistent reasoning.

4.2 Probabilities over Possible Worlds

An obvious way to generate an uncertainty function for a finite set of premises is to assign probabilities to the set of interpretations of the premises and then sum the probabilities over all interpretations that *fail* to support a given premise. There is both a *decision* and an *optimisation* version of the problem. Since we are primarily interested in the analysis of inference from inconsistent premises, for our purpose the *optimisation* version is of more interest.

Before we formally define our problem, we'll first fix some notations. Given a set $\Gamma = \{A_1, \ldots, A_m\}$ with n distinct variables, the set of interpretations over Γ, written as $\mathcal{W}_\Gamma$, is the set of truth assignments over Γ restricted to variables occurring in Γ. Since Γ is assumed to have n distinct variables, $|\mathcal{W}_\Gamma|$ is exactly 2^n. We call each $v_i \in \mathcal{W}_\Gamma$ a possible world for Γ and assume that there is an arbitrary but fixed enumeration of these possible worlds. For each $i \leq 2^n$, we assign a probability $P_i \in [0, 1]$ to v_i. Furthermore we require that $\sum_{i=1}^{2^n} P_i = 1$. Intuitively, P_i is the probability that v_i

is the actual world or the actual outcome. Next we define a $m \times 2^n$ binary matrix $A = (a_{ij})$ by setting

$$a_{ij} = \begin{cases} 0 & \text{if } v_j \models A_i \\ \\ 1 & \text{otherwise} \end{cases} \tag{4.1}$$

We'll call the matrix A an uncertainty matrix for Γ. We note that in the standard literature, a_{ij} is set to 1 if $v_j \models A_i$ and 0 otherwise. But since we take *uncertainty* to be probability of A_i to fail, it is convenient for us to use (4.1) to define $U(A_i)$ instead.

We note that different enumerations of Γ and $\mathcal{W}_\Gamma$ would generate different uncertainty matrices $A = (a_{ij})$ with different orderings on their rows and columns. But we'll treat these matrices as belonging to the same equivalence class invariant under the usual row and column rotation. We define the uncertainty of A_i by setting

$$U(A_i) = \sum_{j=1}^{2^n} a_{ij} P_j = \pi_i \tag{4.2}$$

The set of equalities and inequalities can be written concisely in the matrix notation:

$$\begin{aligned} \mathbf{1}\,\overline{P} &= 1 \\ A\,\overline{P} &= \overline{\pi} \\ \overline{P} &\geq 0 \end{aligned} \tag{4.3}$$

where $\mathbf{1}$ is a 2^n unit row vector, $\overline{P}$ and $\overline{\pi}$ are the column vectors $[P_1, \ldots, P_{2^n}]^T$ and $[\pi_1, \ldots, \pi_m]^T$ respectively. As usual we use $[\ldots]^T$ to denote the *transpose* of $[\ldots]$. We call each distinct column vector $\overline{P}$ a probability distribution for Γ and each distinct column vector $\overline{\pi}$ an uncertainty vector for Γ. We note that the uncertainty of A_i expressed in (4.2) is defined relative to a given probability distribution $\overline{P}$. Where it is necessary, we'll use '$U_{\overline{P}}(A_i)$' to denote the uncertainty of A_i relative to $\overline{P}$. We note that relative to any finite Γ and probability distribution $\overline{P}$, the function $U_{\overline{P}}(\ )$ satisfies the usual Kolmogorov's Axioms.

The decision version of our problem (uncertainty satisfiability problem, **USAT**) can now be stated as follows: given an uncertainty vector $\overline{\pi}$ for Γ, is there a probability distribution $\overline{P}$ for Γ such that the set of equalities and inequalities in (4.3) holds with respect to $\overline{P}$ and $\overline{\pi}$? We note that although there are uncountably many probability distributions for Γ and thus the search space is infinite, our problem is effectively decidable. Any of the standard algorithms (e.g. various versions of the simplex method or the Fourier-Motzkin elimination procedure [82] p. 36–39) for solving linear systems can be used as a decision procedure. If the answer to our decision problem is 'yes', then we say that $\overline{\pi}$ is an *uncertainty assignment* for Γ.

Now consider the addition of one more sentence B with unknown uncertainty π_{m+1}. In the *optimisation* version, we wish to minimise (or maximise) the value $U(B) = \pi_{m+1}$ subject to the constraints imposed by (4.3). Again in matrix notion, we have:

$$\begin{aligned} \text{min (max) } \pi_{m+1} &= \sum_{j=1}^{2^n} a_{m+1,j} P_j \\ \text{subject to} & \\ \mathbf{1}\,\overline{P} &= 1 \\ A\,\overline{P} &= \overline{\pi} \\ \overline{P} &\geq 0 \end{aligned} \tag{4.4}$$

We note that in (4.4), A is now a $(m+1) \times 2^n$ matrix where n is now the number of variables occurring in both Γ and B. Clearly, the optimisation version of **USAT** is just a linear programming problem and thus can be effectively solved using the simplex method. In general, linear programs of interest are those in which the number of unknowns exceeds the number of equations. So for our purpose, we'll assume that in (4.4) $m + 1 \leq 2^n$. Since a set of premises $\{A_1, \ldots, A_m\}$ may not be *logically independent* in the sense that some premise may be provable from the remaining premises, we make no assumption about the *rank* of the matrix A in (4.4).

4.3 Bounded USAT and Inconsistencies

The decision and optimisation versions of **USAT** expressed in (4.3) and (4.4) should be distinguished from a further version of **USAT**. In the *bounded* version of **USAT**, (4.4) is replaced with the inequalities:

$$A\,\overline{P} \leq \overline{\pi} \tag{4.5}$$

We call any such $\overline{\pi}$ in (4.5) which has a solution a bound vector. A vector with identical entries in all of its coordinate is a uniform vector. In the more general version of bounded **USAT**, $A\,\overline{P}$ is bounded from both above and below:

$$\underline{\pi} \leq A\,\overline{P} \leq \overline{\pi} \tag{4.6}$$

The decision version of bounded **USAT** is of particular interest since it captures the classical notion of logical consistency and inconsistency. In particular we have the follow equivalences:

Theorem 4.3.1
Let $\Gamma = \{A_1, \ldots, A_m\}$ *and let* A *be a uncertainty matrix of* Γ*. The following statements are equivalent:*

1. *Every vector* $\overline{\pi} \in [0,1]^m$ *is a bound vector for* Γ*, i.e. for every vector* $\overline{\pi} \in [0,1]^m$*,* $A\,\overline{P} \leq \overline{\pi}$ *hold for some* $\overline{P}$*.*
2. *Every uniform vector* $\overline{\pi} \in [0,1]^m$ *is a bound vector for* Γ*, i.e. for every uniform vector* $\overline{\pi} \in [0,1]^m$*,* $A\,\overline{P} \leq \overline{\pi}$ *hold for some* $\overline{P}$*.*
3. Γ *is classically consistent.*

Proof:
(1. $\Rightarrow$ 2.): Trivial.

(2. $\Rightarrow$ 3.): If Γ is inconsistent, then every column of A must contain at least one entry of 1. Thus for every probability distribution $\overline{P}$ of Γ, there must be an $i \leq m$ such that $U(A_i) = \pi_i > 0$. Hence $A\,\overline{P} \not\leq [0, \cdots\cdots, 0]^T$ for every $\overline{P}$.

(3. $\Rightarrow$ 1.): If Γ is consistent, then there must be a j such that the j-th column of A is the vector $[0, \cdots\cdots, 0]^T$. Let $\overline{P}$ be the probability distribution where $P_j = 1$. Clearly, $A\,\overline{P} = [0, \cdots\cdots, 0]^T$ and so $A\,\overline{P} \leq \overline{\pi}$ for every uncertainty vector $\overline{\pi}$. ∎

The implication of the equivalence of (2) and (3) is that for any inconsistent set Γ, there would be uniform vectors that are not bound vectors for Γ. Of those uniform vectors that are bound vectors for an inconsistent Γ, the most interesting one is of course the minimal one, i.e. one that satisfies inequality (4.5), but no (co-ordinate wise) smaller vector also satisfies (4.5). For minimally inconsistent sets, there is a straightforward way to determine a minimal uniform bound. But first we note that for an inconsistent set, the sum of the uncertainties of its members is bounded below by 1:

Theorem 4.3.2
If Γ is inconsistent then, for any probability distribution $\overline{P}$, $\sum_{A\in\Gamma} U(A) \geq 1$.

Proof:
Let n be the number of variables occurring in Γ. Let A be an uncertainty matrix for Γ. Consider an arbitrary probability distribution $\overline{P} = [P_1, \ldots P_{2^n}]^T$ and equalities:

$$A\,[P_1, \ldots, P_{2^n}]^T \;=\; [\pi_1, \ldots, \pi_{|\Gamma|}]^T \tag{4.7}$$

Clearly it follows from (4.7) that

$$\pi_1 + \ldots + \pi_{|\Gamma|} \;=\; \sum_{i=1}^{|\Gamma|} \sum_{j=1}^{2^n} a_{ij} P_j \tag{4.8}$$

Since Γ is inconsistent, each column of A must contain at least one entry of 1. Thus for each $P_j \in \overline{P}$, there must be an $i \leq |\Gamma|$ such that $a_{ij} = 1$ and thus $1 \times P_j$ must occur at least once in the RHS of (4.8). But $\overline{P}$ is a probability distribution and thus $\sum_{j=1}^{2^n} P_j = 1$. This implies in particular that RHS of (4.8) ≥ 1. We conclude that $\sum_{A\in\Gamma} U(A) \geq 1$. ∎

The converse of theorem (4.3.2) is obviously also true since the consistency of Γ must be witnessed by a zero column in A so distributing maximum probability into such a column yields $\sum_{A\in\Gamma} U(A) = 0$ immediately.

Theorem 4.3.3
Let Γ be a minimally inconsistent set of formulae such that $|\Gamma| = m$. Then the minimal uniform bound vector for Γ is exactly $\overline{m^{-1}}$ (= $[\underbrace{m^{-1} \ldots m^{-1}}_{m\ times}]^T$).

Proof:
Suppose Γ is minimally inconsistent and $|\Gamma| = m$. Let n be the number of variables occuring in Γ. Let A be an uncertainty matrix for Γ. By the minimal inconsistency of Γ, A must contain the identity matrix of order m as a sub-matrix. Without loss of generality we may assume that the identity sub-matrix occupies the left most position of A, otherwise we may perform the usual row and column operations to put A into such a configuration. Let the probability distribution

$$\overline{D} = [\overbrace{\underbrace{m^{-1}, \ldots, m^{-1}}_{m\ times}, 0, \ldots, 0}^{2^n\ times}]^T$$

Clearly we have $A\overline{D} = \overline{m^{-1}}$, so $\overline{m^{-1}}$ is a uniform bound vector for Γ. It's minimality follows from theorem (4.3.2) since we have $\sum \overline{m^{-1}} = m(m^{-1}) = 1$. ∎

In addition, the size of the largest minimal inconsistent subset of Γ gives an absolute lower bound on the number of variables in Γ:

Theorem 4.3.4
If $\Gamma = \{A_1, \ldots, A_m\}$ is minimally inconsistent, then there are at least $\lceil \log_2 m \rceil$ many distinct variables occurring in Γ.

Proof:
We assume that $|\Gamma| = m$ and Γ is minimally inconsistent. Towards a contradiction suppose that there are k variables occurring in Γ where $k < \log_2 m$. Since there are k variables occurring in Γ, there are exactly 2^k distinct valuations for Γ. By the minimal inconsistency of Γ however, for each $i \leq m$, there must be a distinct valuation v_i satisfying $\Gamma \backslash \{A_i\}$. Hence there must be at least m distinct valuations. But by the initial assumption $k < \log_2 m$ and so $2^k < m$. This contradicts the minimal inconsistency of Γ. Hence $k \geq \log_2 m$. But $k \in \mathbb{Z}^+$, hence $k \geq \lceil \log_2 m \rceil$. ∎

Corollary 4.3.1
If m is the cardinality of the largest minimal inconsistent subset of Γ, then there are at least $\lceil \log_2 m \rceil$ many distinct variables occurring in Γ.

4.4 Geometric Rendering of Inconsistencies

According to theorem (4.3.3), the minimal uniform uncertainty bound of a minimal inconsistent set is inversely proportional to the size of the set. Thus the larger the set, the smaller the bound. The geometric relationship between the uniform bound vectors and the uncertainty vectors (of a minimally inconsistent Γ) can be displayed easily in the 2-dimensional and 3-dimensional cases. The dimension here simply corresponds to the cardinality of the set Γ.

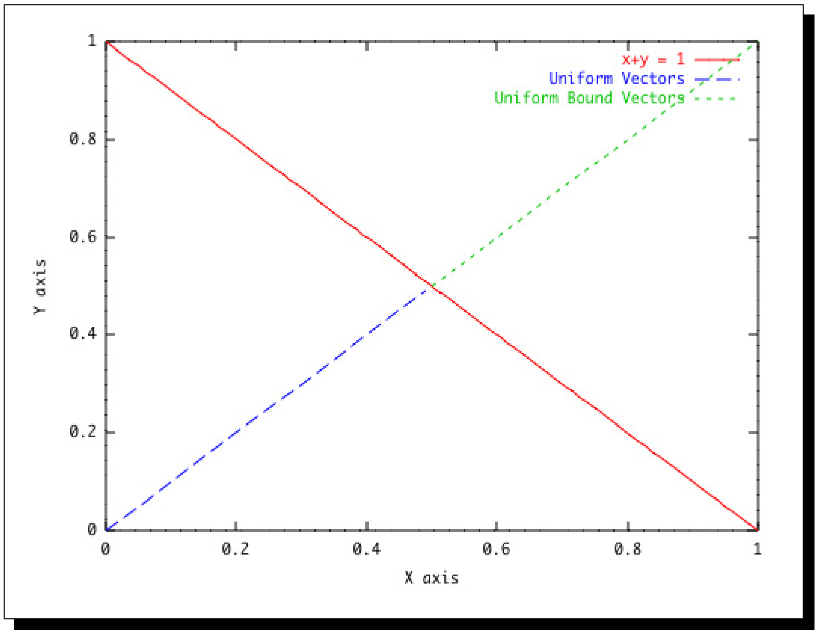

Figure 4.1: Uniform Bounds and Uncertainties in 2D

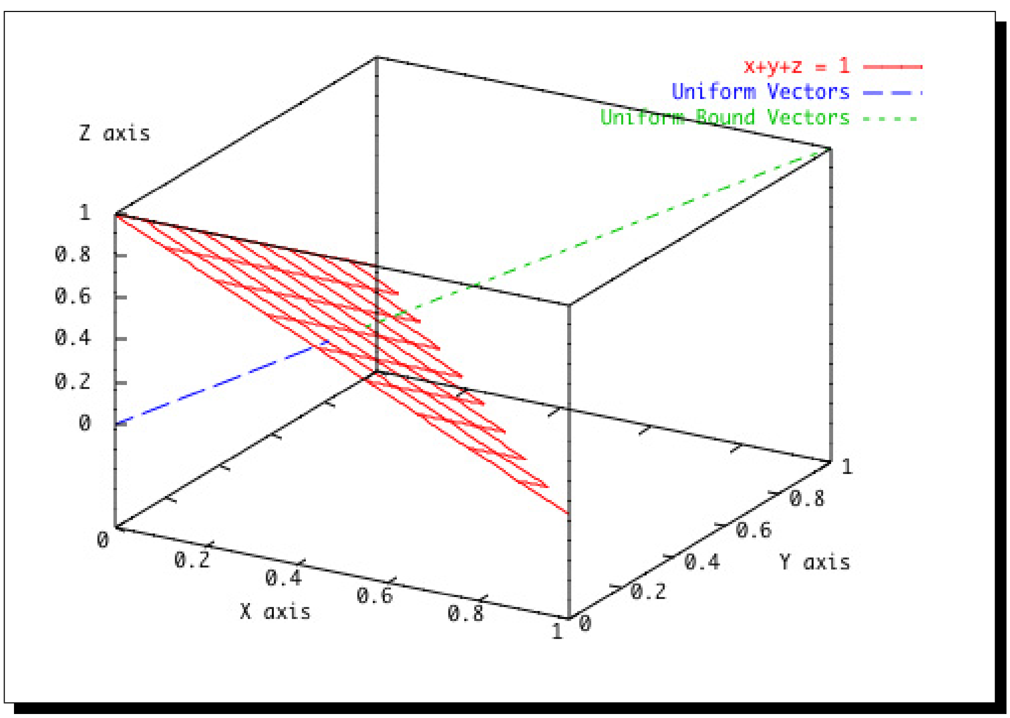

Figure 4.2: Uniform Bounds and Uncertainties in 3D

In the 2-dimensional case (figure (4.1)), the set of all possible uncertainty assignment of Γ are contained within the area of the unit square on or above the diagonal $x + y = 1$. Recall that in the proof of theorem (4.3.3), the sum of all the co-ordinates of an uncertainty assignment is ≥ 1. The set of uniform vectors form the other main diagonal of the unit square. The minimal bound vector is simply the intersection of the two diagonals.

In the 3-dimensional case (figure (4.2)), the set of all possible uncertainty assignments for Γ is contained within the unit cube on or above the plane $x + y + z = 1$. The set of uniform vectors again forms the main diagonal joining the origin and its opposing vertex. The minimal bound vector in this case is simply the intersection of the diagonal and the plane.

The generalisation is obvious. For any minimal inconsistent set of size m, the set of uncertainty assignments is contained within the unit hyper-cube on or above the hyper-plane $\sum_{i=1}^{m} x_i = 1$. The set of uniform vectors forms the diagonal from the origin to the opposing vertex $(1, \ldots, 1)$. The minimal bound vector is again the intersection of the diagonal and the hyper-plane.

Theorem 4.4.1
For any minimal inconsistent set Γ, the minimal bound vector for Γ is the minimal uniform bound vector for Γ.

Proof:
Recall that the euclidean distance between $\overline{x}, \overline{y} \in \mathbb{R}^m$ is defined by

$$e(\overline{x}, \overline{y}) = \left| \sqrt{\sum_{i=1}^{m} (x_i - y_i)^2} \right| \tag{4.9}$$

Thus the euclidean distance between any $\overline{x} = \langle x_1, \ldots x_m \rangle \in [0, 1]^m$ and the origin is $\left| \sqrt{\sum_{i=1}^{m} x_i^2} \right|$. From theorem (4.3.2), we know that for any inconsistent $\Gamma = \{A_1, \ldots, A_m\}$, any uncertainty assignment $\overline{x} \in [0, 1]^m$ for Γ is such that $\sum_{i=1}^{m} x_i \geq 1$. There are two cases to consider:

(case 1) $\sum_{i=1}^{m} x_i = 1$: We want to find value $\overline{x} \in [0, 1]^m$ with minimal distance to the origin subject to the constraint that $\sum_{i=1}^{m} x_i = 1$, i.e.

$$\text{minimise} \quad e(\overline{x}, \mathbf{0}) = \left| \sqrt{\sum_{i=1}^{m} x_i^2} \right| \tag{4.10}$$

$$\text{subject to} \quad \sum_{i=1}^{m} x_i = 1 \tag{4.11}$$

Without loss of generality we may consider minimising the square of (4.10) instead, i.e.

$$f(\overline{x}) = \sum_{i=1}^{m} x_i^2 \tag{4.12}$$

We note that the constraint (4.11) is a closed and bounded subset of the hyper-cube on which f is continuous, thus an absolute minimum value must occur. To find the minima we let $g(\overline{x}) = \sum_{i=1}^{m} x_i$ and use Lagrange multipliers on all partial derivatives of f and g:

$$\nabla f(\overline{x}) = \left(\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_m}\right) = (2x_1, \ldots, 2x_m)$$

whereas

$$\nabla g(\overline{x}) = \left(\frac{\partial g}{\partial x_1}, \ldots, \frac{\partial g}{\partial x_m}\right) = (1, \ldots, 1)$$

Solving for the Lagrange multiplier λ in $\nabla f(\overline{x}) = \lambda \nabla g(\overline{x})$ yields:

$$\begin{aligned} \forall i, 1 \leq i \leq m, \ 2x_i &= \lambda \\ &\implies x_1 = x_2 = \ldots = x_m \end{aligned}$$

But $(x_1, \ldots, x_m)$ is on the hyper-plane $\sum_{i=1}^{m} x_i = 1$. Hence for each i, $x_i = m^{-1}$. The other remaining possible locations for extrema to occur are the endpoints or the point $\overline{x}$ where $\nabla g(\overline{x}) = \mathbf{0}$. We note however that for no $\overline{x} \in [0, 1]^m$ do we have $\nabla g(\overline{x}) = \mathbf{0}$. Hence the latter case is impossible after all. Now each of the endpoints of g is of the form $\overline{x_e} = (\ldots, 0, 1, 0, \ldots)$. Hence for each endpoint $\overline{x_e}$, $f(\overline{x_e}) > f(\overline{m^{-1}})$. Hence we conclude that the $\overline{m^{-1}}$ must be the absolute minimum.

(case 2) $\sum_{i=1}^{m} x_i > 1$: since $[0, 1]^m$ is a euclidean space, if $\overline{x} = (x_1, \ldots, x_m)$ is such that

$\sum_{i=1}^{m} x_i > 1$, then there must be a $\overline{y} = (y_1, \ldots, y_m)$ such that $\sum_{i=1}^{m} y_i = 1$ and

$$e(\overline{x}, \mathbf{0}) \quad = \quad e(\overline{y}, \mathbf{0}) + e(\overline{x}, \overline{y}) \tag{4.13}$$

where $e(\overline{x}, \overline{y}) > 0$. From the previous case however, we know that the absolute minima of f on the hyper-plane $\sum_{i=1}^{m} x_i = 1$ occurs at $\overline{m^{-1}}$. It follows then that $e(\overline{y}, \mathbf{0}) \geq \sqrt{m^{-1}}$. Hence $e(\overline{x}, \mathbf{0}) > \sqrt{m^{-1}}$. Since $\overline{x}$ was arbitrary, we conclude that any such $\overline{x}$ would have $e(\overline{x}, \mathbf{0}) > \sqrt{m^{-1}}$. This suffices to show that $\overline{m^{-1}}$ is the closest uncertainty assignment to the origin. ∎

4.5 Multiple Inconsistencies

The general situation for finding uniform bound vectors for an inconsistent set is considerably more difficult. Although the set of uncertainty assignments is on or above the hyper-plane $\sum_{i=1}^{m} x_i = 1$, we have no information on where the uniform vector may intersect with the set of uncertainty assignments (if they intersect at all). Even if the two do intersect, we have no guarantee that result analogous to theorem (4.4.1) should hold.

In the general case, an inconsistent set may have multiple minimally inconsistent subsets. We say that a set of formulae is contradiction free if it contains no singleton inconsistency. We can obtain some bounds, though not necessarily minimal ones, by looking at the smallest minimal inconsistent subset(s).

Theorem 4.5.1
Let Γ be inconsistent but contradiction free. Let n be the number of variables occurring in Γ and m be the size of the smallest minimal inconsistent subset of Γ. Then there exist a probability distribution $\overline{P}$ such that for all $A \in \Gamma$, $U(A) \leq \frac{2^n - (m-1)}{2^n}$.

Proof:
We let $\Delta = \{A_1, \ldots, A_m\} \subseteq \Gamma$ be a smallest minimal inconsistent subset. Consider the uncertainty matrix A for Γ where the first m row of A correspond to members of Δ. We note that by the minimal inconsistency of Δ, A must be configurable with an identity submatrix of order m in the top left most position, i.e.

$$A = \begin{bmatrix} I_m & B \\ C & D \end{bmatrix}$$

In the worst case B may contain only 1's and thus the maximum possible number of 1's in any given row of the submatrix $[\mathbf{I}_m\ B]$ is $2^n - (m-1)$. We let $\{A_{m+1}, \ldots, A_{|\Gamma|}\} \subseteq \Gamma$ be the set of formulae corresponding to the rows of the submatrix $[C\ D]$. Let A_j be an arbitrary but fixed element of $\{A_{m+1}, \ldots, A_{|\Gamma|}\}$ and consider

$$\Pi = \left\{\Sigma \cup \{A_j\} :\ \Sigma \subset \Delta \text{ and } |\Sigma| = (m-2)\right\}$$

We note that given $|\Delta| = m$, we have $|\Pi| = \binom{m}{m-2}$. Since every $\Sigma \cup \{A_j\} \in \Pi$ is of size $(m-1)$, there must be a $v_\Sigma \in \mathcal{W}_\Gamma$ which witnesses the consistency of $\Sigma \cup \{A_j\}$.

Claim 1: For no $v \in \mathcal{W}_\Gamma$ do we have v witnesses the consistency of more than 2 members of Π.

Proof of Claim 1: Suppose to the contrary that there is some $v \in \mathcal{W}_\Gamma$ which witnesses the consistency of some distinct $\Sigma_1 \cup \{A_j\}, \Sigma_2 \cup \{A_j\}, \Sigma_3 \cup \{A_j\} \in \Pi$. It follows that v must witness the consistency of $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \{A_j\}$. But note that for each $i \in \{1, 2, 3\}$, $|\Sigma_i| = m - 2$ and $\Sigma_i \subset \Delta$. So given that Σ_1, Σ_2 and Σ_3 are all distinct it follows that $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 = \Delta$ and thus $\Delta \subseteq \Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \{A_j\}$. Given that Δ is inconsistent, $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \{A_j\}$ must be inconsistent. This contradicts the initial assumption that v witnesses the consistency of $\Sigma_1 \cup \Sigma_2 \cup \Sigma_3 \cup \{A_j\}$ and hence for no $v \in \mathcal{W}_\Gamma$ do we have v being the witness of more then 2 members of Π.

Claim 2: There are at least $(m-1)$ entries of zero's in each row of the submatrix $[C\ D]$.

Proof of Claim 2: We consider 3 cases:

Case 1: $m = 2$. We note that since Γ contains no contradictions, each row of A must contain at least one entry of 0.

Case 2: $m = 3$. Then $|\Pi| = \binom{m}{m-2} = \binom{3}{1} = 3$. But by claim (1) no single $v \in \mathcal{W}_\Gamma$ can witness the consistency of all three members of Π. Hence to witness each member of Π requires at least two distinct $u, v \in \mathcal{W}$ It follows that every row of $[C\ D]$ must contain at least two entries of 0.

Case 3: $m \geq 4$. Clearly given claim (1), at least $\lceil \frac{|\Pi|}{2} \rceil$ many distinct $v \in \mathcal{W}_\Gamma$ are required to witness the consistency of each member of Π. But note that for $m \geq 4$ we have

$$\begin{aligned} \left\lceil \frac{\binom{m}{m-2}}{2} \right\rceil &\geq \frac{1}{2}\binom{m}{m-2} \\ &= \frac{1}{2} \times \frac{m \times (m-1)}{2} \\ &\geq (m-1) \end{aligned}$$

Hence there are at least $(m-1)$ entries of 0 in each row of $[C\ D]$ as required.

Given claim (2) the maximum possible number of 1's in each row of $[C\ D]$ must be $2^n - (m-1)$. Hence,

$$A[2^{-n}, \ldots, 2^{-n}]^T \leq [\frac{2^n - (m-1)}{2^n}, \ldots, \frac{2^n - (m-1)}{2^n}]^T$$

∎

We note that the bound given in theorem (4.5.1) is an absolute bound. However it is not an attractive bound since 2^n grows exponentially with n and thus the bound approaches 1 very quickly as n grows. One obvious way to obtain a lower bound is to consider the size of Γ instead.

Theorem 4.5.2
Let Γ be inconsistent but contradiction free. Let n be the number of variables occurring in Γ and m be the size of the smallest minimal inconsistent subset of Γ. Then there exists a probability distribution $\overline{P}$ such that for all $A \in \Gamma$, $U(A) \leq \frac{|\Gamma|-(m-1)}{|\Gamma|}$.

Proof:
Let $|\Gamma| = k$. Since m is the size of the smallest inconsistent subset of Γ, every subset $\Delta \subseteq \Gamma$ of size $m-1$ must be consistent. The number of such subsets is $\binom{k}{m-1}$. Let these subsets be enumerated as $\Delta_1, \ldots, \Delta_{\binom{k}{m-1}}$.
For each subset Δ_i, $1 \leq i \leq \binom{k}{m-1}$, there must be a $\nu \in \mathcal{W}_\Gamma$ that witnesses the consistency of Δ_i. For each Δ_i choose one such witness ν_{Δ_i} and set

$$t_i(\nu) = \begin{cases} \binom{k}{m-1}^{-1} & \text{if } \nu = \nu_{\Delta_i} \\ 0 & \text{otherwise} \end{cases} \tag{4.14}$$

We note that the v_{Δ_i}'s are not necessarily unique, i.e. for $i \neq j$, $v_{\Delta_i} = v_{\Delta_j}$ but Δ_i may be distinct from Δ_j.

For each j, $1 \leq j \leq 2^n$, we define the probability:

$$P_j \quad = \quad t_1(v_j) + \ldots + t_{\binom{k}{m-1}}(v_j)$$

We let the probability distribution $\overline{P} = [P_1, \ldots P_{2^n}]^T$, i.e.

$$\overline{P} \quad = \quad \begin{bmatrix} t_1(v_1) + \ldots + t_{\binom{k}{m-1}}(v_1) \\ t_1(v_2) + \ldots + t_{\binom{k}{m-1}}(v_2) \\ \vdots \\ t_1(v_{2^n}) + \ldots + t_{\binom{k}{m-1}}(v_{2^n}) \end{bmatrix} \tag{4.15}$$

Claim 1: $\overline{P}$ is a probability distribution over $\mathcal{W}_\Gamma$, i.e. $\sum_{j=1}^{2^n} P_j = 1$.

Proof of claim 1: We note that for each i, $1 \leq i \leq \binom{k}{m-1}$,

$$\sum_{j=1}^{2^n} t_i(v_j) \quad = \quad \binom{k}{m-1}^{-1}$$

Hence

$$\sum_{i=1}^{\binom{k}{m-1}} \sum_{j=1}^{2^n} t_i(v_j) \quad = \quad \binom{k}{m-1} \times \binom{k}{m-1}^{-1} = 1$$

This completes the proof of our claim. We note that in (4.15) there are exactly $\binom{k}{m-1}$ many non-zero terms.

Let A be an arbitrary but fixed member of Γ, we note that there are exactly $\binom{k-1}{m-2}$ many subsets Δ_i containing A. Without loss of generality we may assume that A is contained in the first $\binom{k-1}{m-2}$ subsets Δ_i. We let $[a_1, \ldots, a_{2^n}]$ be the row vector in the uncertainty matrix A for the corresponding $A \in \Gamma$. Clearly we have

$$\begin{aligned}
U(A) &= a_1 \times \left[t_1(v_1) + \ldots + t_{\binom{k}{m-1}}(v_1)\right] \\
&\quad + \ldots \\
&\quad + a_{2^n} \times \left[t_1(v_{2^n}) + \ldots + t_{\binom{k}{m-1}}(v_{2^n})\right] \\
&= a_1 \times t_1(v_1) + \ldots + a_1 \times t_{\binom{k}{m-1}}(v_1) \\
&\quad + \ldots \\
&\quad + a_{2^n} \times t_1(v_{2^n}) + \ldots + a_{2^n} \times t_{\binom{k}{m-1}}(v_{2^n}) \\
&= \begin{pmatrix} a_1 \times t_1(v_1) \\ + \\ \vdots \\ + \\ a_{2^n} \times t_1(v_{2^n}) \end{pmatrix} + \ldots + \begin{pmatrix} a_1 \times t_{\binom{k}{m-1}}(v_1) \\ + \\ \vdots \\ + \\ a_{2^n} \times t_{\binom{k}{m-1}}(v_{2^n}) \end{pmatrix} \qquad (4.16)
\end{aligned}$$

Claim 2: For each i, $1 \le i \le \binom{k-1}{m-2}$, and each j, $1 \le j \le 2^n$, if $t_i(v_j) \neq 0$, then $a_j \times t_i(v_j) = 0$.

Proof of claim 2:

$$\begin{aligned}
t_i(v_j) \neq 0 &\implies v_j \models \bigwedge \Delta_i \\
&\implies v_j \models A \\
&\implies a_j = 0 \\
&\implies a_j \times t_i(v_j) = 0
\end{aligned}$$

It follows from claim (2) and the definition of t_i that there are at least $\binom{k-1}{m-2}$ many zero terms in (4.16). Thus the maximum number of non-zero terms in (4.16) is $\binom{k}{m-1} -$

$\binom{k-1}{m-2}$. But since each non-zero term in (4.16) is equal to $\binom{k}{m-1}^{-1}$, we have

$$U(A) \leq \frac{\binom{k}{m-1} - \binom{k-1}{m-2}}{\binom{k}{m-1}} = 1 - \frac{\binom{k-1}{m-2}}{\binom{k}{m-1}}$$

$$= 1 - \left[\frac{(k-1)!}{[(k-1)-(m-2)]!(m-2)!} \times \frac{[(k-(m-1)]!(m-1)!}{k!}\right]$$

$$= 1 - \left[\frac{1}{[(k-1)-(m-2)]!} \times \frac{[(k-(m-1)]!(m-1)}{k}\right]$$

$$= 1 - \left[\frac{1}{[(k-m+1)]!} \times \frac{(k-m+1)!(m-1)}{k}\right]$$

$$= \frac{k-(m-1)}{k}$$

Since A was arbitrary, we conclude that $A\,\overline{P} \leq \overline{\frac{k-(m-1)}{k}}$. ∎

Since $\frac{k-(m-1)}{k}$ approaches 1 at a rate that is only linearly related to increases in k, theorem (4.5.2) is an improvement over theorem (4.5.1). We note that in obtaining the bound in theorem (4.5.2) we make no assumption about whether members of Γ are independent. Further improvement can be made if we consider only certain subsets of Γ that are independent in a certain sense.

Definition 4.5.1
We say that a set of formulae Γ *is pairwise independent iff for any* $A, B \in \Gamma$*, neither* $A \vdash B$*, nor* $B \vdash A$*. A subset* $\Pi \subseteq \Gamma$ *is said to be a cover of* Γ *iff*

$$\bigcup_{B \in \Pi} \mathbf{Cn}(B) = \bigcup_{A \in \Gamma} \mathbf{Cn}(A)$$

where **Cn** *is the usual closure under classical deduction.*

The notion of a pairwise independent set is an obvious generalisation of the usual notion of independence – a set of formulae is independent if no member of the set is a consequence of the remaining members of the set. Note that if a set Γ is independent in the ordinary sense, then no proper subset of Γ can be inconsistent (though Γ may be minimally inconsistent). Generalising this, sets that are pairwise independent must be contradiction free, whereas sets that are n-independent in the sense that no member of the set is a consequence of any subset of size $n-1$ must be free of any inconsistent subset of size $\leq n-1$. Thus an n-independent inconsistent set must only have minimal inconsistent subsets of size $\geq n$. It is straightforward to verify from definition (4.5.1) that every set of (contradiction free) formulae Γ must contain a pairwise independent cover of Γ.

Theorem 4.5.3

Let Γ be inconsistent but contradiction free. Let m be the size of the smallest minimal inconsistent subset of Γ and Π be any pairwise independent cover of Γ. Then there exists a probability distribution $\overline{P}$ such that for all $A \in \Gamma$, $U(A) \leq \frac{|\Pi|-(m-1)}{|\Pi|}$.

Proof:

Let Γ and Π fulfil the hypotheses.

Claim 1: m is the size of the smallest minimal inconsistent subset of Π.

Proof of Claim 1: Clearly given that $\Pi \subseteq \Gamma$, the size of the smallest minimal inconsistent subset of Π cannot be less than m. We now show that there is a minimal inconsistent subset of Π of size m. Let Δ be a minimal inconsistent subset of Γ with $|\Delta| = m$. By the covering property of Π, for each $A_i \in \Delta$ there must be a $B_k \in \Pi$ such that $B_k \vdash A_i$. We note that for no two distinct $A_i, A_j \in \Delta$ do we have $B_k \vdash A_i$ and $B_k \vdash A_j$ for some $B_k \in \Pi$. For otherwise, $(\Delta \setminus \{A_i, A_j\}) \cup \{B_k\}$ is a minimal inconsistent subset of size $< m$. Let $\Delta' \subseteq \Pi$ be a set with the property that each $A_i \in \Delta$ is implied by exactly one $B_j \in \Delta'$. Clearly, Δ' must be minimally inconsistent and of size m. This completes the proof of our claim.

Claim 2: For any formulae A, B if $A \vdash B$ then for all probability distributions $\overline{P}$, $U(B) \leq U(A)$.

Proof of Claim 2: Suppose that $A \vdash B$. Consider an arbitrary but fixed $l \times 2^n$ uncertainty matrix A with row $r_i = [a_{i1}, \ldots, a_{i2^n}]$ corresponding to A and row $r_j = [b_{j1}, \ldots, b_{j2^n}]$ corresponding to B. Clearly for any $k \leq 2^n$, if $a_{ik} = 0$ then the corresponding $b_{jk} = 0$. So for any arbitrary probability distribution $\overline{Q} = [Q_1, \ldots Q_{2^n}]^T$, we have

$$\sum_{k=1}^{2^n} b_{jk} \times Q_k \quad \leq \quad \sum_{k=1}^{2^n} a_{ik} \times Q_k$$

But since A and $\overline{Q}$ were completely arbitrary, we conclude that $U(B) \leq U(A)$ on any probability distribution.

By the covering property of Π every $A \in \Gamma \setminus \Pi$ is implied by some $B \in \Pi$, so it follows from claim (2) that for any probability distribution, for each $A \in \Gamma \setminus \Pi$ there exists some $B \in \Pi$ such that $U(A) \leq U(B)$. By theorem (4.5.2) and claim (1) however, there must be a probability distribution $\overline{P}$ such that for all $B \in \Pi$, $U(B) \leq \frac{|\Pi|-(m-1)}{|\Pi|}$. If $\overline{P}$ is not defined for all of Γ, it is trivial to extend $\overline{P}'$ for all of Γ such that for any $B \in \Pi$, $U(B) \leq \frac{|\Pi|-(m-1)}{|\Pi|}$ still holds with respect to $\overline{P}'$. But then claim (2) confirms that on $\overline{P}'$, $U(A) \leq \frac{|\Pi|-(m-1)}{|\Pi|}$ holds for any $A \in \Gamma \setminus \Pi$. Hence we conclude that on $\overline{P}'$, $U(A) \leq \frac{|\Pi|-(m-1)}{|\Pi|}$ holds for any $A \in \Gamma$. ∎

The bound obtained in theorem (4.5.3) clearly improves as the value of m approaches $|\Pi|$. This shows that in a large data set, the uncertainty bound of any single datum is better for dispersed inconsistencies than for concentrated inconsistencies.

We also note that in the event that Γ is minimally inconsistent, Γ must be a pairwise independent cover of itself. Thus applying theorem (4.5.3) to Γ we get the uncertainty bound $|\Gamma|^{-1}$ which is in accordance with theorem (4.3.3). In light of this, theorem (4.5.3) can be taken to be a generalisation of theorem (4.3.3).

4.6 Uncertain Inference

As we have already noted, premises that are inconsistent (but contradiction free) have non-trivial uncertainties. In this section we would like to continue the investigation initiated by Adams and Levine in [5] and examine how uncertainties may be transmitted from premises to conclusions. In [1; 2; 3], Adams extended the result of [5] to cover a language with a conditional connective. However the issue of uncertainties transmitted from *inconsistent* premises to conclusions has not been addressed in any of their subsequent works. We start by identifying several candidate (uncertainty) entailment relations – all of which can be said to preserve the uncertainty bound of the premises under some sense:

Definition 4.6.1
For any set of formulae Γ *and formula* B *we define the following entailment between* Γ *and* B

Certainty Entailment: *For any probability distribution* $\overline{P}$ *such that* $U_{\overline{P}}(A) = 0$ *for all* $A \in \Gamma$, *we have* $U_{\overline{P}}(B) = 0$. *We denote this by* $\Gamma \models_0 B$.

Uncertainty Entailment: *For any probability distribution* $\overline{P}$ *such that* $U_{\overline{P}}(A) < 1$ *for all* $A \in \Gamma$*, we have* $U_{\overline{P}}(B) < 1$*. We denote this by* $\Gamma \models_{<1} B$.

ϵ-Entailment: *For any* $\epsilon \in [0, 1]$*, for any probability distribution* $\overline{P}$ *such that* $U_{\overline{P}}(A) \leq \epsilon$ *for all* $A \in \Gamma$ *we have* $U_{\overline{P}}(B) \leq \epsilon$*. We denote this by* $\Gamma \models_{\leq\epsilon} B$.

We note that an inconsistent set of formulae cannot all be certain together. Thus certainty entailment is an explosive entailment for any inconsistent premises. In fact it is just classical entailment:

Proposition 4.6.1
Certainty entailment is equivalent to classical entailment.

Proof:
We note that any $v \in \mathcal{W}_\Gamma$ which verifies all of Γ but falsifies B would also confirm the existence of a $\overline{P}$ with all $A \in \Gamma$ having $U(A) = 0$ but $U(B) \neq 0$.

Conversely if $\Gamma \models B$, then for the uncertainty matrix A of $\Gamma \cup \{B\}$, every column $[a_{1j}, \ldots a_{|\Gamma|+1j}]^T$ with all 0's in the first $|\Gamma|$ entries (corresponding to members of Γ) will have $a_{|\Gamma|+1j} = 0$ (corresponding to B). Let $\overline{P} = [P_1 \ldots P_{2^n}]^T$ be an arbitrary but fixed probability distribution such that for all $A_k \in \Gamma$, $U_{\overline{P}}(A_k) = 0$. Clearly for each $A_k \in \Gamma$

$$U_{\overline{P}}(A_k) = \sum_{j=1}^{2^n} a_{kj} P_j$$

is zero if either $a_{kj} = 0$ or $P_j = 0$ for each $i \leq 2^n$. If $P_j = 0$ then obviously $a_{|\Gamma|=1j} \times P_j = 0$. But if $P_j \neq 0$ then $a_{kj} = 0$ for each $k \leq |\Gamma|$. But then $a_{|\Gamma|=1j} = 0$ as well and thus $a_{|\Gamma|=1j} \times P_j = 0$. Hence $\sum_{j=1}^{2^n} a_{|\Gamma|+1j} P_j = 0$, i.e. $U_{\overline{P}}(B) = 0$. ∎

Turning now to uncertainty entailment, it is clearly an improvement over certainty entailment for handling inconsistencies. The basic idea of uncertainty entailment is that if each of the $A_i \in \Gamma$ is free from complete uncertainty, then the conclusion B is also free of complete uncertainty. Since contradiction free inconsistent premises have non-trivial uncertainties, the antecedent of the conditional in our definition is never falsified in such a case. Thus we do not have $A, \neg A \models_{<1} B$ in general. But note that in the presence of contradictions, we do have $A \wedge \neg A \models_{<1} B$. Moreover for any classical tautology $\top$ we have $\models_{<1} \top$ trivially. In fact the logic which captures $\models_{<1}$ completely is the *discursive* logic(s) developed by Jaśkowski in [97]. For a complete sequent formulation of *discursive* logic, the reader can consult the system S of Knight in [111]. But the basic idea of a *discursive* logic is to take the union of all the theorems of an underlying logic $\vdash_L$ together with the deductive closures (under $\vdash_L$) of each singleton of the premises, i.e. for any Γ and B we have

$$\Gamma \vdash_D B \text{ iff } \vdash_L B \text{ or } A \vdash_L B \text{ for some } A \in \Gamma$$

For a different choice of the underlying L we get a different discursive logic.

Proposition 4.6.2
For any Γ and B, $\Gamma \models_{<1} B$ iff either $\vdash B$ or $A \vdash B$ for some $A \in \Gamma$ where $\vdash$ is the usual classical propositional logic.

Proof:
As noted before it is trivially true that if B is a classical tautology, then for any probability distribution $\overline{P}$ we have $U_{\overline{P}}(B) \leq U_{\overline{P}}(A)$ for any A. So we'll consider any B that is not a tautology. For the only if direction we assume that $A \vdash B$ for some $A \in \Gamma$. Then from claim (2) of theorem (4.5.3), for every probability distribution $\overline{P}$ over $\Gamma \cup \{B\}$, we have $U_{\overline{P}}(B) \leq U_{\overline{P}}(A)$. So in particular for any probability distribution $\overline{Q}$ with $U_{\overline{Q}}(C) < 1$ for every $C \in \Gamma$ we have $U_{\overline{Q}}(B) < 1$. This shows that $\Gamma \models_{<1} B$.

For the if direction, we assume that for no $A \in \Gamma$ do we have $A \vdash B$. Now consider the uncertainty matrix $A = (a_{ij})$ for $\Gamma \cup \{B\}$. As usual we'll assume that A is a $(m+1) \times 2^n$ matrix, where the first m rows correspond to members of Γ and the $(m+1)$-th row corresponds to B. Moreover we assume that t is the number of 1's occurring in the $(m+1)$-th row of A. We note that $t \geq 1$ since B is not a tautology by the initial assumption. We define the probability distribution $\overline{P}$ as follows: for every j, $1 \leq j \leq 2^n$,

$$P_j = \begin{cases} t^{-1} & \text{if } a_{m+1,j} = 1 \\ \\ 0 & \text{otherwise} \end{cases}$$

Clearly given how $\overline{P}$ is defined, $U_{\overline{P}}(B) = [\sum_{j=1}^{2^n}(a_{m+1,j} \times P_j)] = 1$. But note that since for each $A \in \Gamma$, $A \not\vdash B$ so we have, for each $i \leq m$ there must be a $j_i \leq 2^n$ such that $a_{ij_i} = 0$ but $a_{m+1,j_i} = 1$. This implies that for each $A \in \Gamma$ we have $U_{\overline{P}}(A) \leq \frac{t-1}{t} < 1$. Thus $\overline{P}$ witnesses the failure of $\Gamma \models_{<1} B$. ∎

Turning now to ϵ-entailment, the basic requirement is that the uncertainty of a conclusion B should never exceed the maximum value of the uncertainty of any given $A \in \Gamma$, i.e. for any probability distribution $\overline{P}$, $U(B) \leq \max\{U(A) : A \in \Gamma\}$. As it turns out ϵ-entailment is in fact equivalent to uncertainty entailment:

Proposition 4.6.3

For any Γ *and* B, $\Gamma \models_{\leq \epsilon} B$ *iff* $\Gamma \models_{<1} B$.

Proof:

The only if direction is trivial since $\models_{<1}$ is a special case of $\models_{\leq \epsilon}$ when $\epsilon < 1$.

For the if direction, consider B where B is a tautology. Then for any Γ we have $\Gamma \models_{\leq \epsilon} B$ since $U(B) = 0$ for any probability distribution. Suppose then that B is not a tautology but for some $A \in \Gamma$, $A \vdash B$ holds. Then from claim (2) of theorem (4.5.3) again, for every probability distribution $\overline{P}$ over $\Gamma \cup \{B\}$, we have $U(B) \leq U(A)$. Thus we have $U(B) \leq \max\{U(A) : A \in \Gamma\}$ as required. ∎

To put the matter in terms of preservation, discursive logic is exactly the logic which preserves the uncertainty bounds of premises. Note however that discursive logic does not allow for full *aggregation* of premises. In general we have $U(\bigwedge_{i=1}^{m} A_i) \leq \sum_{i=1}^{m} U(A_i)$, but not $U(\bigwedge_{i=1}^{m} A_i) \leq \max\{U(A_i) | 1 \leq i \leq m\}$. In light of this, discursive logic is a very extreme approach to bounding the uncertainty of the conclusion. When the value of $\max\{U(A) : A \in \Gamma\}$ is close to 1, it is of course desirable to ensure that the conclusion's uncertainty should not exceed this bound. But when the value of $\max\{U(A) : A \in \Gamma\}$ is small, a slightly riskier inference with a higher conclusion uncertainty may be acceptable. More importantly, aggregation is particularly useful for fusing information from multiple sources. We'll introduce a kind of entailment relation which permits a limited form of aggregation by bounding the size of the aggregating set. Our entailment relation also provides a partial solution to a problem stated in Knight [111] (page 360). But first we need to fix some terminologies and definitions.

Definition 4.6.2

Let $k \in \mathbb{Z}^+$ *be arbitrary but fixed. Let* Γ *be a finite set of formulae in* n *variables. The set of all subsets of* Γ *of size* $\leq k$ *is denoted by* $\wp_k(\Gamma)$.

If $\overline{P} = [P_1, \ldots, P_{2^n}]^T$ *is a probability distribution over* Γ, *we say that* $\overline{P}$ *is* i*-positive if* $P_i > 0$.

If $\Delta \subseteq \Gamma$, *we say that* $\overline{P}$ *verifies* Δ *if there exists an* $i \leq 2^n$ *such that* $\overline{P}$ *is* i*-positive and the* i^{th} *term of* $U_{\overline{P}}(A)$ *is* 0 *for each* $A \in \Delta$, *i.e. where* A *is the uncertainty matrix for* Γ

and $j(1), \ldots, j(t)$ *are the respective enumeration of members of* Δ*, we have* $a_{j(1)i} \times P_i = \ldots = a_{j(t)i} \times P_i = 0$ *under* $\overline{P}$.

Note that if $\Delta \in \wp_k(\Gamma)$ is inconsistent, then no $\overline{P}$ will verify Δ. Intuitively, $\overline{P}$ verifies a Δ only if $\overline{P}$ distributes non-zero probability into at least one model of Δ. We now introduce a generalised version of ϵ-entailment with an additional parameter k as a bound on the size of the aggregating set.

Definition 4.6.3
Let $k \in \mathbb{Z}^+$ *be arbitrary but fixed. Let* Γ *be any finite set of formulae in* n *variables. For any formula* B*, we say that* Γ k*-entails* B*,* $\Gamma \models_k B$*, iff* $U(B) < 1$ *on every probability distribution* $\overline{P}$ *which verifies every* $\Delta \in \wp_k(\Gamma)$*.*

Now for different choices of k we can regain different degrees of aggregation. So for instance if $k \geq 2$ and $A_1, A_2 \in \Gamma$ then $\Gamma \models_k A_1 \wedge A_2$. Again k is the absolute upper bound on the number of (independent) members of Γ that can be conjoined. Note also that any tautology $\top$ is k-entailed by any Γ since $U(\top) = 0 < 1$ holds trivially. Moreover if the size of the smallest minimal inconsistent subset of Γ is m and $m < k$, then no $\overline{P}$ will verify every $\Delta \in \wp_k(\Gamma)$ and thus $\Gamma \models_k B$ for any B holds trivially, i.e. $\models_k$ explodes when $m < k$. We summarise the properties of $\models_k$ in theorem (4.6.1). The content of our theorem is self-explanatory. Part (1) shows that $\models_k$ is an extension of $\models_{\leq\epsilon}$. Part (2) shows that $\models_k$ is a kind of substructural logic. Part (3) shows that $\models_k$ is monotonically increasing with respect to k. Part (4) is a generalised version of proposition (4.6.2) and therein shows that $\models_k$ can be viewed as a kind of generalised discursive logic (and thus is decidable). Part (5) shows that $\models_k$, like $\models_{\leq\epsilon}$, preserves the uncertainty bound of the premises in a certain sense.

Theorem 4.6.1

1. *For any* $k \in \mathbb{Z}^+$*,* $\models_{\leq\epsilon} \subseteq \models_k$*.*
2. $\models_k$ *is reflexive and monotonic but transitivity fails.*
3. *If* $k', k \in \mathbb{Z}^+$ *and* $k' < k$*, then* $\models_{k'} \subseteq \models_k$*.*
4. *Let* $k \in \mathbb{Z}^+$ *be fixed. For any* Γ *and* B*,* $\Gamma \models_k B$ *iff* $B \in \bigcup\{\mathbf{Cn}(\Delta) | \Delta \in \wp_k(\Gamma)\}$*.*
5. *Let* $\epsilon \in [0, 1]$ *such that* $\epsilon < 1$*. Let* $\Gamma \models_k B$*. Then for any probability distribution* $\overline{P}$*, if* $\sum_{A \in \Delta} U_{\overline{P}}(A) \leq \epsilon$ *holds for each* $\Delta \in \wp_k(\Gamma)$ *then* $U_{\overline{P}}(B) \leq \epsilon$*.*

Proof:
(1): As noted before for any tautology B, $\Gamma \models_k B$ holds trivially. So we'll assume that B

is not a tautology and for an arbitrary Γ, we have $\Gamma \models_{\leq \epsilon} B$. From propositions (4.6.2) and (4.6.3), it follows that for some $A \in \Gamma$, $A \vdash B$. Again from claim (2) of theorem (4.5.3) it follows that for any probability distribution $\overline{Q}$, we have $U_{\overline{Q}}(B) \leq U_{\overline{Q}}(A)$. Clearly $\{A\} \in \wp_k(\Gamma)$ for any $k \in \mathbb{Z}^+$. So if $\overline{P}$ verifies every $\Delta \in \wp_k(\Gamma)$, it must also verify $\{A\}$. This implies the existence of some i such that $P_i > 0$ and the i^{th}-term of $U_{\overline{P}}(A)$ is 0. Since $A \vdash B$, the i^{th} term of $U_{\overline{P}}(B)$ must be 0 as well. Thus $U_{\overline{P}}(B) \leq [(\sum_j^{2^n} P_j) - P_i] < 1$.

(2): For reflexivity, clearly if $A \in \Gamma$ then $\{A\} \in \wp_k(\Gamma)$ for any $k \in \mathbb{Z}^+$. So if $\overline{P}$ verifies every $\Delta \in \wp_k(\Gamma)$, it must also verify $\{A\}$ as well. This implies that for some i, $U(A) \leq [(\sum_j^{2^n} P_j) - P_i] < 1$ as required.

For monotonicity, we note that $\wp_k(\Gamma) \subseteq \wp_k(\Gamma, \Sigma)$ so if $\overline{P}$ verifies every member $\wp_k(\Gamma, \Sigma)$, it must also verify every member of $\wp_k(\Gamma)$. So on the assumption that $\Gamma \models_k A$ holds $\Gamma, \Sigma \models_k A$ must hold as well.

To see the failure of transitivity, consider $\Gamma = \{p,\ \neg p \vee r,\ \neg r\}$. We have

$$\Gamma \models_2 p \wedge (\neg p \vee r) \text{ and } \Gamma, p \wedge (\neg p \vee r) \models_2 q$$

But note that $\Gamma \not\models_2 q$.

(3) We note if $k' < k$ then $\wp_{k'}(\Gamma) \subseteq \wp_k(\Gamma)$ for any Γ. Thus if $\overline{P}$ verifies every member of $\wp_k(\Gamma)$ it must also verify every member of $\wp_{k'}(\Gamma)$. So on the assumption that $\Gamma \models_{k'} A$, $\Gamma \models_k A$ must hold as well.

(4) For the if direction let $\overline{P}$ be any probability distribution which verifies every $\Delta \in \wp_k(\Gamma)$. We want to show that $U_{\overline{P}}(B) < 1$ on the assumption that $B \in \bigcup\{\mathbf{Cn}(\Delta) |\ \Delta \in \wp_k(\Gamma)\}$. So we assume that for some $\Delta_0 \in \wp_k(\Gamma)$, $\Delta_0 \vdash B$. By the initial assumption however $\overline{P}$ must verify Δ_0, so there exists some i such that $P_i > 0$ and the i^{th} term of $U_{\overline{P}}(A)$ is 0 for every $A \in \Delta_0$. But $\Delta_0 \vdash B$ so the i^{th} term of $U_{\overline{P}}(B)$ is 0 as well. As in (1) and (2), this suffices to show that $U_{\overline{P}}(B) \leq [(\sum_{j=1}^{2^n} P_j) - P_i] < 1$.

For the only if direction, we assume that $B \notin \bigcup\{\mathbf{Cn}(\Delta) |\ \Delta \in \wp_k(\Gamma)\}$, i.e. for every $\Delta \in \wp_k(\Gamma)$, $\Delta \not\vdash B$. We'll show the existence of a $\overline{P}$ which verifies every $\Delta \in \wp_k(\Gamma)$ but on $\overline{P}$ we have $U_{\overline{P}}(B) = 1$.

Consider the uncertainty matrix $A = (a_{ij})$ for $\Gamma \cup \{B\}$. As usual we'll assume that A is a $m + 1 \times 2^n$ matrix, where the first m rows correspond to members of Γ and the

$(m+1)$-th row corresponds to B. Moreover we assume that t is the number of 1's occurring in the $(m+1)$-th row of A. We note that $t \geq 1$ since B cannot be a tautology by the initial assumption. We define the probability distribution $\overline{P}$ as follows: for every j, $1 \leq j \leq 2^n$,

$$P_j = \begin{cases} t^{-1} & \text{if } a_{m+1,j} = 1 \\ \\ 0 & \text{otherwise} \end{cases}$$

Clearly given how $\overline{P}$ is defined, $U_{\overline{P}}(B) = [\sum_{j=1}^{2^n}(a_{m+1,j} \times P_j)] = 1$.

Claim: $\overline{P}$ verifies every $\Delta \in \wp_k(\Gamma)$.

Proof of claim: Let $\Delta \in \wp_k(\Gamma)$ be arbitrary. By the initial assumption $\Delta \nvdash B$ so there must be a column in A which witnesses this. Let the witnessing column be the s^{th} column in A. We note that $\overline{P}$ must be s-positive since $P_s = t^{-1} > 0$. Moreover the s^{th} term of $U(A)$ must be 0 for every $A \in \Delta$. Hence $\overline{P}$ verifies Δ. Since Δ was arbitrary, this suffices to show that $\overline{P}$ verifies every member of $\wp_k(\Gamma)$.

(5) We assume that $\Gamma \models_k B$ and that $\overline{P} = [P_1 \ldots P_{2^n}]^T$ is an arbitrary probability distribution such that $\sum_{A \in \Delta} U_{\overline{P}}(A) \leq \epsilon < 1$ holds for each $\Delta \in \wp_k(\Gamma)$. From (4) above it follows that $B \in \bigcup\{\mathbf{Cn}(\Delta) | \ \Delta \in \wp_k(\Gamma)\}$. This implies that for some $\Delta_0 \in \wp_k(\Gamma)$ we have $\Delta_0 \vdash B$. But by the initial assumption $\sum_{A \in \Delta_0} U_{\overline{P}}(A) \leq \epsilon < 1$. By theorem (4.3.2)), it follows that Δ_0 must be consistent. Let $|\Gamma| = m$ with n variables and let $|\Delta_0| = t$. Without loss of generality we may assume that the first t rows of the uncertainty matrix A correspond to members of Δ_0 and the $(m+1)$-th row of A corresponds to B. Using the usual column rotation, A can be reconfigured into the following sub-matrices:

$$A = \begin{bmatrix} B & C \\ D & E \end{bmatrix}$$

B is a $t \times s$ submatrix with each column containing at least one entry of 1; C is a $t \times (2^n - s)$ zero submatrix. By the consistency of Δ_0, C cannot be empty. We note that since $\Delta_0 \vdash B$, the last row of E must be 0's. This gives the following absolute upper bound on $U_{\overline{P}}(B)$:

$$U_{\overline{P}}(B) \quad \leq \quad P_1 + \ldots + P_t$$

However we note that since each column of B contains at least one entry of 1, we have the following absolute lower bound on $\sum_{A \in \Delta_0} U_{\overline{p}}(A)$:

$$P_1 + \ldots + P_t \quad \leq \quad \sum_{A \in \Delta_0} U_{\overline{p}}(A)$$

Hence $U_{\overline{p}}(B) \leq \sum_{A \in \Delta_0} U_{\overline{p}}(A) \leq \epsilon < 1$ as required. ∎

4.7 Bounded Reasoning in Natural Deduction

Although our motivation for $\models_k$ has been stated solely in probabilistic terms so far, we should point out that $\models_k$ can also be regarded as a kind of *resource bounded reasoning*. If we take the suggestion of the linear logician seriously and treat premises as resources to be consumed, it is natural to be concerned with how premises are used and propagated in a proof. In certain natural deduction systems for classical logics, formulae can be labelled with numerals to facilitate bookkeeping of premise dependence as well as to keep track of the subproof structure of a given proof. Lemmon's classic text [118] for instance uses such a labelling device. Anderson and Belnap also introduce labels to a Fitch style natural deduction formulation of relevant logics in their seminal work [7].[1] The basic function of the labels is to serve as names for the premises and the process of deduction involves propagating the labels from premises to the conclusion in a systematic and controlled way. Of course in classical logics, a premise can be reused as often as required in a proof. In linear logics however, this is no longer the case – we may have $A \vdash_L B$ but not $A, A \vdash_L B$. Hence in linear logics, the fundamental data structure of premises is *multisets* instead of the usual sets. Of course k-entailment is not a linear logic and does not require any accounting for how many times a given premise is used in a proof. But $\models_k$ does require a mechanism to keep track of how many distinct premises (from Γ) are used in a given proof. But we must be careful to distinguish between different ways of introducing assumptions into a proof. Using Lemmon's system in [118] as a point of reference, there are at least 4 distinct ways to introduce assumptions into a proof:

1. the rule of assumption introduction (AI)

[1] The use of labels was first introduced by Jaśkowski in [96] and subsequently refined by Quine [141] and Suppes [175]. For a history and discussion of various versions of natural deduction systems, see Pelletier [137]. For a completely general approach to logics via the use of Labelled Deductive Systems, see Gabbay [73].

2. $\rightarrow$-introduction (CP for conditional proof),
3. $\vee$-elimination rule, ($\vee - E$), and
4. the reductio ad absurdum rule (RAA)

The rule that is of immediate interest to us is the rule of AI:

		$\vdots$	
$\{\alpha\}$	(i)	A	Assumption
		$\vdots$	

Figure 4.3: Assumption Introduction

The notation is slightly modified here with the label α enclosed in set brackets. The usual convention is that the label α is just a numeral indicating that the introduced assumption A depends on itself. We'll continue to use our convention from now on. The rule of AI allows us to introduce an assumption at any stage in a proof and the rule can be used for any number of assumptions in a given proof. The basic idea is that given a set of premises Γ, we can introduce finitely many members of Γ into a proof via the use of AI. Clearly AI cannot be a valid rule for $\models_k$. Nonetheless, it is possible to port AI into $\models_k$ by using double labelling – both keeping track of the assumption dependence and keeping count of the number of assumptions introduced from the given set Γ. For an arbitrary but fixed k, we let the set of labels for the rule of assumption introduction be $\mathbb{N} \times \{1, \ldots, k\}$:

The modified AI rule allows us to introduce an assumption into a proof at line (i) provided that for no $j < i$ do we have $\langle j, \beta \rangle$ ($\beta \in \{1, \ldots k\}$) occurring as a label for a distinct assumption introduced by AI.

		$\vdots$	
$\{\langle i, \beta \rangle\}$	(i)	A	Assumption Intro
		$\vdots$	

Figure 4.4: Modified Assumption Introduction

The net effect of our modification is that no more than k assumptions from a given Γ can be introduced into a proof. What about other rules which also require the introduction of assumptions? The simplest approach to take is to deploy a distinct set of labels for these rules. We take the set of labels for assumptions introduced via CP, $\vee$-E and RAA to be $\mathbb{N} \times \{0\}$. CP can be redefined using our labelling system as follows:

		$\vdots$	
$\{\langle i, 0\rangle\}$	(i)	A	Assumption CP
		$\vdots$	
$\mathcal{J}$	(j)	B	...
$\mathcal{J} \setminus \{\langle i, 0\rangle\}$	$(j+1)$	$A \to B$	$i - j$ CP
		$\vdots$	

Figure 4.5: Modified rule of $\to$-Introduction

		$\vdots$	
$\mathcal{I}$	(i)	$A \vee B$	
		$\vdots$	
$\{\langle j, 0\rangle\}$	(j)	A	Assumption $\vee - E$
		$\vdots$	
$\mathcal{L} \setminus \{\langle j, 0\rangle\}$	(l)	C	...
		$\vdots$	
$\{\langle m, 0\rangle\}$	(m)	B	Assumption $\vee - E$
		$\vdots$	
$\mathcal{N} \setminus \{\langle m, 0\rangle\}$	(n)	C	...
$(\mathcal{I} \cup \mathcal{L} \cup \mathcal{N}) \setminus \{\langle j, 0\rangle, \langle m, 0\rangle\}$	(n+1)	C	$i, j - l, m - n \vee - E$
		$\vdots$	

Figure 4.6: Modified $\vee - E$

Note that at line $(j+1)$ (figure (4.5), the label $\langle i, 0\rangle$ is removed from the set of labels $\mathcal{J}$ and thereby *discharging* (an occurrence of) the assumption A at line (i). The occurrences of formulae corresponding to labels in $\mathcal{J} \setminus \{\langle i, 0\rangle\}$ are said to be *undischarged* at line $(j + 1)$. Note that the notion of a discharged and undischarged assumption is defined over occurences of an assumption. An occurence of an assumption A may be discharged at line k while a distinct occurrence of A may be undischarged at k. The dotted line at line (j) represents some rule in the system which is being applied at line (j) of the proof. The block beginning with the line with the assumption A and ending with the line with A being discharged is a subproof of the overall proof. Note also that although A is discharged at line $(j+1)$, the label $\langle i, 0\rangle$ may not be a member of $\mathcal{J}$ at all. The $\rightarrow$ introduced by the CP rule is the material implication – it allows us to obtain $A \rightarrow B$ trivially if we can obtain B without actually using the assumption A.

In figure (4.6), the inferred statement C at line $(n+1)$ inherits all undischarged assumptions of $A \vee B$ at line (i), as well as those from line (l) and line (n). Both A and B (at line (j) and line (m) respectively) are discharged at line $(n+1)$. Also note that there are two subproof structures involved here – blocks $(j-l)$ and $(m-n)$. For RAA the assumption A is discharged at line $(j+1)$ as usual.

		$\vdots$	
$\{\langle i, 0\rangle\}$	(i)	A	Assumption RAA
		$\vdots$	
$\mathcal{J}$	(j)	$B \wedge \neg B$	...
$\mathcal{J} \setminus \{\langle i, 0\rangle\}$	$(j+1)$	$\neg A$	$i - j$ RAA
		$\vdots$	

Figure 4.7: Modified RAA

$$\frac{\begin{array}{cc} \mathcal{I} & A \\ & \vdots \\ \mathcal{J} & A \rightarrow B \end{array}}{\begin{array}{cc} \mathcal{I} \cup \mathcal{J} & B \end{array}}\ \text{MP} \qquad \frac{\begin{array}{cc} \mathcal{I} & \neg B \\ & \vdots \\ \mathcal{J} & A \rightarrow B \end{array}}{\begin{array}{cc} \mathcal{I} \cup \mathcal{J} & \neg A \end{array}}\ \text{MT}$$

$$\frac{\begin{array}{cc} \mathcal{I} & A \\ & \vdots \\ \mathcal{J} & B \end{array}}{\begin{array}{cc} \mathcal{I} \cup \mathcal{J} & A \wedge B \end{array}}\ \wedge - \text{I} \qquad \frac{\begin{array}{cc} \mathcal{I} & A \wedge B \end{array}}{\begin{array}{cc} \mathcal{I} & B \end{array}}\ \wedge - \text{E}$$

$$\frac{\begin{array}{cc} \mathcal{I} & A \end{array}}{\begin{array}{cc} \mathcal{I} & A \vee B \end{array}}\ \vee - \text{I} \qquad \frac{\begin{array}{cc} \mathcal{I} & A \end{array}}{\begin{array}{cc} \mathcal{I} & \neg\neg A \end{array}}\ \text{DN}$$

Figure 4.8: Modified Lemmon's Rules

All remaining rules of Lemmon's system involve no introduction of assumptions and can be redefined in our double labelling system. We'll simplify the graphical representations of these rules by omitting some details here. We note that DN is an invertible rule. $\vee - \mathsf{I}$, $\wedge - \mathsf{I}$ and $\wedge - \mathsf{E}$ are all commutative with respect to $\vee$ and $\wedge$.

Note that each application of our rules involves no discharging of assumptions; thus all undischarged assumptions prior to the application of a rule will remain undischarged at the line in which the rule is applied. We'll call all the above rules non-discharging rules and call CP, $\vee - \mathsf{E}$ and RAA discharging rules. We'll identify Lemmon's original system as L' and our modified system as L. The notion of provability in L is defined in the usual way.

Definition 4.7.1
We say that a sentence B is L-provable from a set of sentences Γ, written as $\Gamma \vdash_{\mathsf{L}} \mathsf{B}$, iff there is a finite sequence $\langle\langle\mathcal{I}_1, \mathsf{A}_1\rangle, \ldots, \langle\mathcal{I}_s, \mathsf{A}_s\rangle\rangle$ such that

1. $\mathsf{A}_s = \mathsf{B}$,

2. *$\mathcal{I}_s$ is either empty or contains only labels introduced by AI in the sequence. In the second case, all undischarged occurrences of assumptions are members of Γ.*

3. *For each $t \leq s$, A_t is either an assumption introduced from the set Γ via AI or introduced by CP, $\vee - \mathsf{E}$ or RAA, or obtained from previous line(s) via either a non-discharging or a discharging rule, and the set of labels $\mathcal{I}_t$ is obtained by the application of the corresponding rule.*

Before we show that $\vdash_{\mathsf{L}}$ is indeed adequate for $\models_k$, we need a number of intermediate results. As usual for any Γ, we set $C_{\mathsf{L}}(\Gamma) = \{\mathsf{B} | \Gamma \vdash_{\mathsf{L}} \mathsf{B}\}$.

Theorem 4.7.1
(see [118]) Lemmon's system L' is (strongly) sound and complete with respect to classical semantics.

An immediate corollary of theorem (4.7.1) is that all tautologies (theorems) of classical logic are L' provable from $\emptyset$, i.e. the last line of any such L' proof is $\langle\emptyset, \top\rangle$ where $\top$ is a tautology.

Lemma 4.7.1
For any tautology $\top$, every L' proof of $\top$ can be converted into an L proof of $\top$ and vice versa.

Proof:
By theorem (4.7.1), there must be a L' proof of $\top$ from $\emptyset$. We note that since CP, $\vee - E$ and RAA are the only discharging rules in L', any L' proof of $\top$ which contains applications of AI can be converted into a L' proof of $\top$ without AI. This holds since any assumption introduced by AI will remain undischarged in the last line of a L' proof if the assumption is actually used in the proof. To see that such a L' proof can be converted into a L proof of $\top$, we observe that the only difference between L and L' is the labelling used. Since AI is not used in such a L' proof of $\top$, each line of the proof can be rewritten with labels from $\mathbb{N} \times \{0\}$. This suffices to show that all tautologies are provable in L from $\emptyset$. To see that the converse also holds, again observe that any L proof of $\top$ can be converted into a AI free L proof of $\top$. Each line of such a L proof can be rewritten with the corresponding label in L'. ∎

Lemma 4.7.2
L *has the deduction theorem, i.e. for any* A, B *and* Γ, $\Gamma, A \vdash_L B$ *only if* $\Gamma \vdash_L A \rightarrow B$.

Proof:
We assume that $\Gamma, A \vdash_L B$ and the proof of B from Γ, A to be

$$D = \langle \langle \mathcal{I}_1, C_1 \rangle, \ldots, \langle \mathcal{I}_s, B \rangle \rangle$$

There are two cases to consider:

(case 1): There is no occurrence of A as an assumption introduced via AI in D. Then D can be extended to a L proof of $A \rightarrow B$ from Γ as follows:

		$\vdots$	
$\mathcal{I}_s$	(s)	B	...
$\{\langle s+1, 0 \rangle\}$	$(s+1)$	A	Assumption CP
$\mathcal{I}_s \cup \{\langle s+1, 0 \rangle\}$	$(s+2)$	$A \wedge B$	$s, s+1, \wedge - I$
$\mathcal{I}_s \cup \{\langle s+1, 0 \rangle\}$	$(s+3)$	B	$s+2, \wedge - E$
$\mathcal{I}_s$	$(s+4)$	$A \rightarrow B$	$(s+1) - (s+3)$, CP

We note that by the initial assumption, $\mathcal{I}_s$ contains no label corresponding to any occurrence of A. Hence at line $(s+4)$ the only undischarged assumptions are members of Γ.

(case 2): There is an occurrence of A in D via the rule of AI. Suppose that A occurs at line (i) with label $\langle i, j \rangle$, with $j \in \{1, \ldots, k\}$. We'll construct an L proof, D' from D as follows: line (i) of D is replaced with the assumption A for CP with the label $\langle i, 0\rangle$. Each subsequent line of D which uses A with label $\langle i, j\rangle$ will be replaced with the label $\langle i, 0\rangle$. At line $(s+1)$, $A \to B$ is obtained via the use of CP and thereby discharges the occurrence of A at line (i):

		$\vdots$	
$\{\langle i, 0\rangle\}$	(i)	A	Assumption CP
		$\vdots$	
$\mathcal{I}'_s$	(s)	B	$\ldots$
$\mathcal{I}'_s \setminus \{\langle i, 0\rangle\}$	$(s+1)$	$A \to B$	$(i-s)$ CP

We note that the set of labels $\mathcal{I}'_s$ is the same as $\mathcal{I}_s$ except that any occurrence of $\langle i, j\rangle$ in $\mathcal{I}_s$ is replaced by $\langle i, 0\rangle$. Thus at line $(s+1)$ all undischarged assumptions are members of Γ. We conclude that in either case we have $\Gamma \vdash_L A \to B$. ∎

Lemma 4.7.3
Let $k \in \mathbb{Z}^+$ *be arbitrary but fixed. For any* Γ,

$$\bigcup_{\Delta \in \wp_k(\Gamma)} \mathbf{Cn}(\Delta) \subseteq C_L(\Gamma)$$

Proof:
We let $k \in \mathbb{Z}^+$ be fixed and Γ be any arbitrary set of formulae. We make the assumption that for some arbitrary $\Delta \in \wp_k(\Gamma)$ and some arbitrary formula B, we have $\Delta \vdash B$. We let $\Delta = \{A_1, \ldots, A_n\}$. By theorem (4.7.1) we have $\Delta \vdash_{L'} B$. Since the deduction theorem also holds with respect to L', we have

$$\emptyset \vdash_{L'} [A_1 \to (A_2 \to (\ldots(A_n \to B)\ldots))]$$

Hence by lemma (4.7.1) we have

$$\emptyset \vdash_L [A_1 \to (A_2 \to (\ldots(A_n \to B)\ldots))$$

We let $D = \langle\langle \mathcal{I}_1, C_1\rangle, \ldots, \langle \mathcal{I}_s, C_s\rangle\rangle$ be the L proof of $C_s = [A_1 \to (A_2 \to (\ldots(A_n \to B)\ldots))$ We note that since $\Delta = \{A_1, \ldots, A_n\} \in \wp_k(\Gamma)$, D can be extended to a proof D' by repeated application of AI and MP:

		$\vdots$	
$\emptyset$	(s)	$[A_1 \to (\ldots(A_n \to B)\ldots]$	$\ldots$
$\{\langle s+1, 1\rangle\}$	$(s+1)$	A_1	Assumption AI
$\{\langle s+2, 2\rangle\}$	$(s+2)$	A_2	Assumption AI
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\{\langle s+n, n\rangle\}$	$(s+n)$	A_n	Assumption AI
$\{\langle s+1, 1\rangle\}$	$(s+n+1)$	$(A_2 \to (\ldots(A_n \to B)\ldots)$	(s), $(s+1)$ MP
$\{\langle s+1, 1\rangle, \langle s+2, 2\rangle\}$	$(s+n+2)$	$(A_3 \to (\ldots(A_n \to B)\ldots)$	$(s+n+1)$, $(s+2)$, MP
$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\{\langle s+1, 1\rangle, \ldots, \langle s+n-1, n-1\rangle\}$	$(s+2n-1)$	$A_n \to B$	$(s+n-1)$, $(s+2n-2)$, MP
$\{\langle s+1, 1\rangle, \ldots, \langle s+n, n\rangle\}$	$(s+2n)$	B	$(s+n)$, $(s+2n-1)$, MP

Note that the undischarged assumptions at line $(s+2n)$ are exactly the members of Δ. It follows that $\Gamma \vdash_L B$ as required. ∎

Lemma 4.7.4

Let $k \in \mathbb{Z}^+$ *be arbitrary but fixed. For any* Γ*,*

$$C_L(\Gamma) \subseteq \bigcup_{\Delta \in \wp_k(\Gamma)} \mathbf{Cn}(\Delta)$$

Proof:

We let $k \in \mathbb{Z}^+$ be arbitrary but fixed. For arbitrary Γ and B we let $\Gamma \vdash_L B$. By lemma(4.7.2), there must be some $\Delta = \{A_1, \ldots, A_n\} \in \wp_k(\Gamma)$ such that

$$\emptyset \vdash_L [A_1 \to (A_2 \to (\ldots(A_n \to B)\ldots))]$$

By lemma (4.7.1) we get

$$\emptyset \vdash_{L'} [A_1 \to (A_2 \to (\ldots(A_n \to B)\ldots))]$$

It follows that $\Delta \vdash_{L'} B$. But $\Delta \in \wp_K(\Gamma)$, hence by theorem (4.7.1) $B \in \bigcup_{\Delta \in \wp_k(\Gamma)} \mathbf{Cn}(\Delta)$ as required. ∎

Theorem 4.7.2

For any arbitrary but fixed $k \in \mathbb{Z}^+$ *we have* $\models_k = \vdash_L$.

Proof:

From part (4) of theorem (4.6.1) we get $\Gamma \models_k B$ iff $B \in \bigcup_{\Delta \in \wp_k(\Gamma)} \mathbf{Cn}(\Delta)$ for arbitrary Γ

and B. So by lemmas (4.7.3) and (4.7.4), we get $\Gamma \models_k B$ iff $\Gamma \vdash_L B$ for arbitrary Γ and B. ∎

4.8 Conclusion

We note that although our uncertainty analyses of inconsistencies and inferences do not seem to provide adequate provisions to deal with contradiction infested premises, we can nonetheless adopt the approach from the previous chapter by rewriting each premise as a set of relevant prime implicates. The result of such a rewrite will undoubtedly affect the uncertainty (bound) of the premises. We have not undertaken any systematic study of the effect on the uncertainty function $U(\ )$ of either restricting or rewriting the syntax of the premises. In the case in which all premises are restricted to clauses, we conjecture that the uncertainty bound can be improved further. More specifically we put forth the following conjecture:

Conjecture 4.8.1
If $\Gamma = \{C_1, \ldots, C_s\}$ *is a set of clauses and* m *is the size of the smallest inconsistent subset of* Γ*, then there is a probability distribution* $\overline{P}$ *such that for all* $i \leq n$*,* $U(C_i) \leq m^{-1}$*. Moreover this is the best possible bound.*

We end this chapter by noting that given our modification of Lemmon's system in section (4.7), we agree with Slaney [171] that nonclassical logics, relevant logics and many of their rivals, are far from being contrived and esoteric. Simple modifications of the labelling procedure can produce logics that are both intuitive and intrinsically interesting.

Chapter 5

QC **Logic**

5.1 Introduction

Logic has long been recognized as the study of reasoning – reasoning not in the psychological sense of how people actually reason or what inferences people tend to draw given some initial assumptions, but reasoning in the sense of providing some standards for evaluating reasoning patterns and distinguishing good ones from bad ones. The development of logic in the past has concentrated on both the proof theoretic and model theoretic aspects of logic. Yet the *pragmatics* aspect of logic seems not to have received the same attention. In this chapter we would like to demonstrate that a logic can be *practical* in the sense that it can assist us in evaluating and measuring the *amount* of information in an inconsistent set of data. Though we envision that any intelligent practical reasoning system must have some mechanism for handling inconsistencies, our goal here is not to address the issue of what is reasonable to conclude given some inconsistent data. Indeed there seems to be no *a priori* reason to favor any one particular inconsistency tolerant system. Rather we would like to illustrate how a particular paraconsistent logic can be used as a tool for analysing inconsistent information. In particular, we would like to be able to quantitatively compare the relative information value of different sets of inconsistent data.

5.2 Paraconsistent Logics

A recalcitrant problem in the development of practical reasoning systems is the issue of uncertainty. One sort of uncertainty is the result of *underdetermination* of data. Another sort is the result of *overdetermination*. All this is well known and is documented in Belnap's [21; 22]. When information gathered from different sources is either incomplete or inconsistent, it is difficult to draw reliable conclusions to guide further action. More importantly, when inconsistencies arise a reasoner must take measures to guard against drawing trivial conclusions. Revising one's data to restore consistency may be an available option, but on occasions it is more important to maintain the integrity of the original data – perhaps the inconsistent data is irrelevant to one's overall project. On other occasions it may even be 'desirable' to have inconsistencies in one's database; for instance, inconsistencies may be deployed as directives to guide learning, and inconsistencies in a taxpayer's records can be used as a reason to prompt further investigation (see [74] for more discussion). The important point is that many ordinary circumstances require us to reason in the presence of inconsistencies. The main motivation for paraconsistent logics is precisely to develop reasoning systems that can tolerate inconsistencies. In classical logic, the rule *ex contradictione quodlibet* is derivable:

$$\frac{A \qquad \neg A}{B}$$

The practical implication of this is that classical logic is unsuitable as a practical reasoning system – it provides no guidance on what can be concluded when inconsistent information is presented, any formula can be derived from an inconsistent set of assumptions. In paraconsistent logics however *ex contradictione quodlibet* is no longer derivable. But as a result paraconsistent logics are also weaker than classical logic. In C. I. Lewis's original proof of *ex contradictione quodlibet* [120], various classical rules are deployed and thus various strategies are open for weakening classical logic:

(1)	$A \wedge \neg A$	Assumption
(2)	A	$1, \wedge - E$
(3)	$\neg A$	$1, \wedge - E$
(4)	$A \vee B$	$2, \vee - I$
(5)	B	$3{,}4 \vee - E$

Figure 5.1: Lewis's Proof of Ex Falso Quodlibet

Ignoring for now the difference between $\{A \wedge \neg A\}$ and $\{A, \neg A\}$, it is clear that we can block the derivation by blocking any one of the rules in line (2), (3), (4) or (5). Relevant logicians, for instance, opt for a solution by blocking the use of $\vee$-E, also known as disjunctive syllogism (see [7; 8]).[1] Logicians favoring analytic implication

[1]We qualify with the warning that relevant logicians do not all agree on this point. See exchanges between Burgess [45; 46; 47], Read [150], Mortensen [130] and Lavers [117].

opt for blocking the use of $\vee$-I, also known as the rule of addition (see [64; 178]). [2] Yet another novel approach is to restrict the order in which the rules are applied. Clearly Lewis's derivation requires that $\vee$-I be used before the use of $\vee$-E. So we can impose restrictions on both $\vee$-I and $\vee$-E so that they cannot be used in that specific combination. The resulting logic is called Quasi-classical logic (QC logic) by Besnard and Hunter in [34] and Hunter in [91]. Indeed a very simple way to characterize QC logic is this: rules in classical logic are divided into *composition* rules and *decomposition* rules; a derivation in QC logic proceeds by first applying decomposition rules and then applying composition rules, but not vice versa.

One of the main advantages of QC logic is that all connectives are interpreted classically as *boolean* connectives. The composition and decomposition rules are divided roughly along the lines of introduction and elimination rules associated with these connectives.[3] Thus we have not changed any of the meanings of $\neg$, $\wedge$ or $\vee$. To simplify matters we take $\neg$, $\wedge$ and $\vee$ to be the primitive connectives and assume that $\wedge$ and $\vee$ are both commutative and associative and satisfy the contraction rules: $\frac{A \vee A}{A}$ $\frac{A \wedge A}{A}$. We take the rules governing the commutativity, associativity and the contraction property of $\wedge$ and $\vee$ to be both decomposition and composition rules, i.e. they can be used at any stage of a QC-proof. The remaining decomposition and composition rules of QC are given in figure (5.2).

A few comments about the rules are in order. The composition rules are, for the most part, the reversal of the decomposition rules. $\neg\neg$-Introduction, C-Distribution and C-de Morgan are the reversal of $\neg\neg$-Elimination, D-Distribution and D-de Morgan respectively. Obviously all our rules are classically valid. But more importantly, all the rules except $\wedge$-Elimination and $\vee$-Introduction preserve exactly the classical models of their premises. By this we mean that any two-valued interpretation is a model of the premises if and only if it is also a model of the conclusion. In the case of $\wedge$-Elimination and $\vee$-Introduction however, the set of models for the premises is properly contained in the set of models for the conclusion, i.e. the conclusions of these rules are strictly weaker then their assumptions. Amongst all the decomposition and composition rules, $\vee$-Introduction is the only rule which allows the introduction of new propositional variables not contained in the premises.

Also note that the set of decomposition rules is sufficient to reduce any formula to its CNF and thus to an equivalent set of clauses. We can further obtain the *resolvants*

[2] See [158] for a detailed discussion of these positions.

[3] We qualify here that strictly speaking disjunctive syllogism is not an elimination rule associated with the connective $\vee$. Note that in stating DS we are required to invoke both $\neg$ and $\vee$.

Decomposition Rules

$\wedge$-Elimination $$\frac{A \wedge B}{A}$$

$\neg\neg$-Elimination $$\frac{\neg\neg A \vee B}{A \vee B} \qquad \frac{\neg\neg A}{A}$$

Resolution $$\frac{A \vee B \qquad \neg A \vee C}{A \vee C} \qquad \frac{A \vee B \qquad \neg A}{B} \qquad \frac{A \qquad \neg A \vee B}{B}$$

D-Distribution $$\frac{A \vee (B \wedge C)}{(A \vee B) \wedge (A \vee C)} \qquad \frac{(A \wedge B) \vee (A \wedge C)}{A \wedge (B \vee C)}$$

D-de Morgan $$\frac{\neg(A \wedge B) \vee C}{(\neg A \vee \neg B) \vee C} \qquad \frac{\neg(A \vee B) \vee C}{(\neg A \wedge \neg B) \vee C}$$

$$\frac{\neg(A \wedge B)}{\neg A \vee \neg B} \qquad \frac{\neg(A \vee B)}{\neg A \wedge \neg B}$$

Composition Rules

$\wedge$-Introduction $$\frac{A \qquad B}{A \wedge B}$$

$\vee$-Introduction $$\frac{A}{A \vee B}$$

$\neg\neg$-Introduction $$\frac{A \vee B}{\neg\neg A \vee B} \qquad \frac{A}{\neg\neg A}$$

C-Distribution $$\frac{(A \vee B) \wedge (A \vee C)}{A \vee (B \wedge C)} \qquad \frac{A \wedge (B \vee C)}{(A \wedge B) \vee (A \wedge C)}$$

C-de Morgan $$\frac{(\neg A \vee \neg B) \vee C}{\neg(A \wedge B) \vee C} \qquad \frac{(\neg A \wedge \neg B) \vee C}{\neg(A \vee B) \vee C}$$

$$\frac{\neg A \vee \neg B}{\neg(A \wedge B)} \qquad \frac{\neg A \wedge \neg B}{\neg(A \vee B)}$$

Figure 5.2: Composition and Decomposition Rules

from these clauses via the use of the resolution rule. Normally the use of the resolution rule in automated theorem proving aims at deriving the empty clause. But in our case, the role of resolution is to decompose clauses into literals so that we can identify and isolate all the inconsistencies in the assumptions.

Unlike a natural deduction system with a linear representation for proofs, a derivation in QC is best represented as a tree-like structure.[4] Officially a derivation in QC logic takes a (finite) set of formulae $\Gamma = \{A_1, \ldots, A_i\}$ as assumptions and a formula B as a conclusion. We write $\Gamma \vdash_{QC} B$ to denote that there is a QC derivation of B from Γ. The derivation proceeds first by the construction of a series of decomposition trees via the decomposition rules. The leaves of these decomposition trees are simply members of Γ; nodes are formulae obtained via the application of the decomposition rules, and finally their roots are either clauses or resolvants of clauses obtained by application of the resolution rule. The roots of the decomposition trees are then used, as leaves, to construct a composition tree via the composition rules. The composition tree terminates when it reaches the conclusion B. The overall structure of a QC proof is given in figure (5.3).

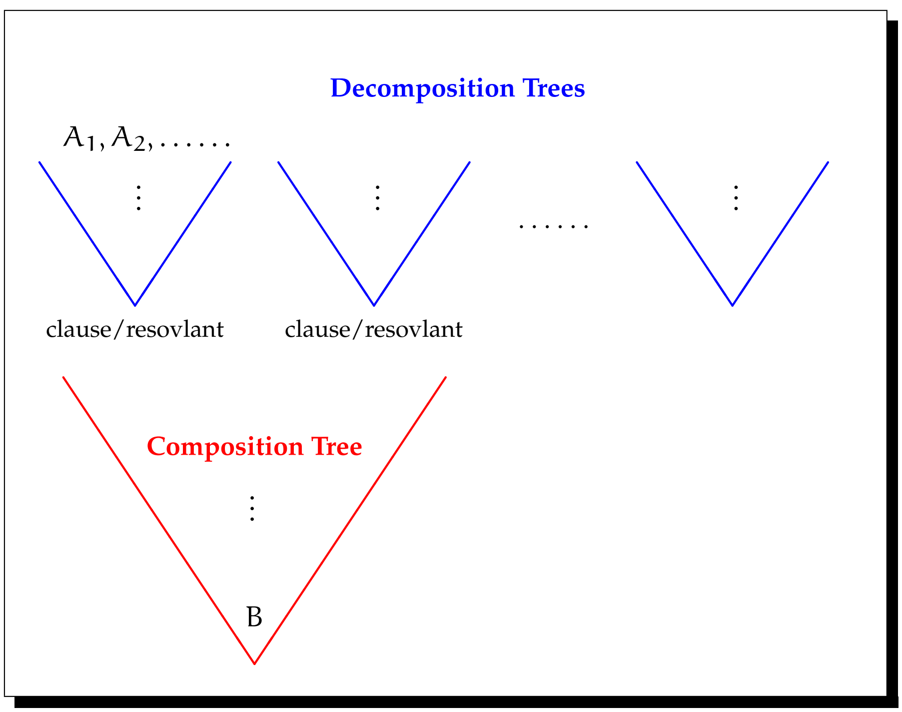

Figure 5.3: The Structure of a QC Proof.

[4]Given the use of double labeling introduced in section (4.7), QC can be ported to a Lemmon style or a Fitch style proof system. Once again we need two disjoint sets of labels – one for the decomposition rules and one for the composition rules.

We should point out that the definition of a QC proof given here is one presented in Hunter [91]. A different (and non-equivalent) definition is given in [34; 35]. In the alternative version, all decomposition and composition rules can be applied in any order except $\vee - I$ can only be used as the last step in a proof. We shall refer to this alternative system as QC′.

The underlying proof theory of QC is reminiscent of the dual tableau systems developed by Rasiowa and Sikorski in [149]. RS systems turn out to be a very flexible framework for a variety of logics and have been further developed by Konikowska and Avron [15; 114; 115] in recent years. In standard tableau systems for classical logics, the validity of a formula B is proven by showing that $\neg B$ has a closed tableau, i.e. every branch contains a contradictory pair of formulae. But to determine whether a branch in a tableau is closed, we need to keep a history of the nodes in the branch. In RS systems however, there is no need to keep such a history. The termination condition for a branch is always determined by the current node on the branch. In RS systems the validity of B is proven by showing that B has a *correct* decomposition tree. The key components of a RS system are

- decomposition rules
- expansion rules
- fundamental nodes (or sequences)

In RS systems, a decomposition rule is of the form:

$$\frac{\Omega}{\Omega_1|\ldots|\Omega_i}$$

An expansion rule is of the form:

$$\frac{\Pi_1|\ldots|\Pi_i}{\Pi}$$

where all the Ω's and Π's are finite sequences of formulae and '$\ldots|\ldots$' is branch splitting. Unlike a standard tableau rule, branch splitting is to be interpreted conjunctively in RS systems.

A proof of the validity of B in a RS system begins with the application of decomposition rules to B until each branch reaches a node that is either *indecomposable* or *fundamental*. Depending on the particular logic in question, the fundamental nodes or sequences are then extended by the expansion rules. Thus by modifying the components of a RS system, we can achieve different controls over deduction and hence obtain different logics. For standard classical logic, the key feature of the decomposition rules is that they are validity preserving in both directions. To prove the validity of B we require that every branch in the finite decomposition tree of B terminates in a node that is a tautology.

Returning to QC however, $\vdash_{QC}$ is not designed as a system for proving theorems. In fact, QC logic has no theorems, i.e. no formula is derivable from the empty set of assumptions. Moreover, like $\vdash_L$ from section (4.7), $\vdash_{QC}$ does not satisfy the usual transitivity or *cut* rule. However, $\vdash_{QC}$ is both reflexive and monotonic, and like classical logic QC is decidable (see Hunter [91] for details). For our purpose here, the most interesting aspect of QC is its decomposition rules. Recall that in section (3.4.1) we have introduced the notions of *prime implicate* and *relevant prime implicate*. As is well known, the resolution proof procedure is complete with respect to prime implicate generation in the sense that if B is a prime implicate of a formula A, then B is a resolvant of CNF(A). Similarly any set of rules that is sufficient to reduce any formula into its equivalent CNF form is complete with respect to RPI generation. Hence the decomposition rules of QC are both PI and RPI complete. Garson [77] and Priest [140] both observe that certain formulations of resolution theorem provers are implicitly paraconsistent in the way they handle inconsistencies. This is true for the case in which the resolution rule is used for PI generation – arbitrary clauses are *not* derivable from an inconsistent set of clauses in general.

Definition 5.2.1
The decomposition closure of a set Γ*, denoted by* $C_D(\Gamma)$ *is the least superset of* Γ *that is closed under the decomposition rules of* QC *(including the contraction rules for* $\wedge$ *and* $\vee$*).*

We note two important facts about C_D: for any Γ the set of propositional variables occurring in Γ is exactly the set of propositional variables occurring in $C_D(\Gamma)$. Moreover, if Γ is finite, then $C_D(\Gamma)$ is also finite. We may say that C_D is a variable and finiteness preserving closure operator.

As usual we say that a CNF of a formula A is reduced if it is a minimal CNF such that all of its propositional variables are variables occurring in A. We say that a reduced CNF of a formula A respects $C_D(\Gamma)$ if all of its clauses can be composed from members of $C_D(\Gamma)$ via the composition rules. Now to determine whether A is QC derivable from a finite Γ is simply a matter of finding a reduced CNF of A that respects $C_D(\Gamma)$. Though there is no unique reduced CNF for a formula A, it is easy to see that one of them would respect $C_D(\Gamma)$ iff all of them would. Since $C_D(\Gamma)$ is finiteness preserving, $\wp(C_D(\Gamma))$ must be finite given that Γ is finite. Hence there are only finitely many ways to generate composition trees from $C_D(\Gamma)$. The checking must terminate eventually.

A key feature of C_D is its ability to identify literals that are involved in an inconsistency. Other paraconsistent logics such as FDE [21] or da Costa's C_ω [53] lack this feature since they lack the resolution rule. Consider for instance,

Example 5.2.1
For $\Gamma = \{p \vee q,\ p \vee \neg q,\ \neg p \wedge r\}$

$$\begin{array}{ll} \Gamma \vdash_{QC} p & \Gamma \vdash_{QC} \neg p \\ \Gamma \vdash_{QC} q & \Gamma \vdash_{QC} \neg q \\ \Gamma \vdash_{QC} r & \Gamma \nvdash_{QC} \neg r \end{array}$$

In our example there is a clear sense in which the variable r is not involved in any inconsistency though it is conjoined with $\neg p$ which is a culprit. However, we should point out that given the consideration in section (3.4), C_D is subjected to the same criticism raised in example (3.4.3) – applications of the D-Distribution rules will result in disjunctive consequences which may be deemed unacceptable. Once again the underlying issue is how much disjunctively redundant information should be tolerated in the presence of inconsistencies.

Example 5.2.2
Let $\Gamma = \{\mathbf{p}, \neg p, q \vee r, \neg r\}$ and $\Delta = \Gamma \cup \{\mathbf{p} \vee \mathbf{r}\}$. We have

$$\begin{array}{ll} r \notin C_D(\Gamma) & \neg r \in C_D(\Gamma) \\ r \in C_D(\Delta) & \neg r \in C_D(\Delta) \end{array}$$

We note that in full QC, we have both $\Gamma \vdash_{QC} p \vee r$ and $\Delta \vdash_{QC} r$ but $\Gamma \nvdash_{QC} r$. Example (5.2.2) indeed demonstrates the failure of the transitivity of deduction in QC. But more importantly, it reinforces a key point for section (3.4) – disjunctively redundant information such as $p \vee r$, when combined with the resolution rules, allows inconsistencies to be propagated amongst premises. Thus in section (3.4.1) we propose to minimise disjunctive redundancies by considering only the relevant prime implicates of a given formula. Indeed, given the *procedural* nature of a QC proof we can explicitly introduce an additional minimisation step in a QC proof by requiring that only relevant prime implicates of premises or resolvants of relevant prime implicates of premises can be used to construct a composition tree. Instead of using QC's decomposition rules to convert each premise into CNF and then prune each CNF formula into RPI's, we could employ the semantic graph method of algorithm (3.4.2) from chapter 3 to generate RPI's. The collection of RPI's can then be given to a conventional resolution prover to generate resolvants. Once again, in consonance with the general methodology of knowledge compilation, RPI generation can be viewed as *off-line* processing while resolution and composition can be viewed as *run time* query-answering.

5.3 Information Measurement

An old idea about information is that there is an inverse relationship between information and possibility. In Shannon-Weaver communication theory [169; 170] this relationship is expressed by the following equation,[5]

$$I(A) \quad = \quad -\log P(A) \tag{5.1}$$

In equation (5.1), $I(A)$ is the *amount of information* or *information value* conveyed by the message A and $P(A)$ is the probability of A occurring. Not surprisingly, the thrust of the idea is that information eliminates possibilities – the more unlikely that A occurs the more informative it is to assert A. In [17] Barwise called this the inverse relation principle and gave an illuminating account of the interdependence of information and possibility. To illustrate, consider our example from section (1.1). Recall that O is located in one out of nine possible locations represented by a 3×3 grid:

	q_1	q_2	q_3
p_1			
p_2			
p_3			

Figure 5.4: A simple logical representation of an object's location.

The set of all possible locations of O can be regarded as a probability space. Furthermore, we may assume that each possible location has an equal probability weight. Using (5.1), we can calculate the information of values $A = p_1$ and $A' = p_1 \wedge \neg q_2$:

$$I(p_1) = -\log\frac{3}{9} = 0.48 \qquad\qquad I(p_1 \wedge \neg q_2) = -\log\frac{2}{9} = 0.65$$

[5] See [105] chapter 2-3 for an overview. For a related approach to *semantic information theory* see Hintikka [84; 85; 86].

Not surprisingly, we have $I(A) < I(A')$. Even at an intuitive level it is clear that A' is more informative since A' provides the additional information $\neg q_2$.

Shannon's information measure expressed by equation (5.1) has a number of important properties, in particular:

$$I(A) \geq 0 \quad \text{for} \quad 0 \leq P(A) \leq 1 \tag{5.2}$$

$$\lim_{P(A)\to 1} I(A) = 0 \tag{5.3}$$

$$I(A) > I(B) \quad \text{for} \quad P(A) < P(B) \tag{5.4}$$

Thus equations (5.2) and (5.3) say that $I(A)$ is a non-negative quantity that approaches 0 as $P(A)$ approaches 1. Equation (5.4) says that information increases with uncertainty. Furthermore if $A_1, \ldots, A_k$ are successive and independent messages with the joint probability $P(A_1, \ldots, A_k) = P(A_1) \times \ldots \times P(A_k)$ then

$$\begin{aligned} I(A_1, \ldots, A_k) = -\log[P(A_1) \times \ldots \times P(A_k)] &= -\sum_{i=1}^{k} \log P(A_i) \\ &= \sum_{i=1}^{k} I(A_i) \end{aligned} \tag{5.5}$$

In [169], Shannon showed the following remarkable theorem:

Theorem 5.3.1
The information measure defined by equation (5.1) is the only function that satisfies all of the properties from (5.2) to (5.5).

5.3.1 Inconsistent Information

Data, encoded as formulae in a logical language, are representations of the state of the world. For a consistent set of data each classical interpretation of the data can be regarded as a possible state of the world. Since a consistent set of formulae in finitely many propositional variables has only finitely many non-equivalent interpretations, we can treat the collection of all possible non-equivalent interpretations as a probability space and assign equal probability weight to each interpretation. Naturally this leads to a definition of information analogous to equation (5.1).

Definition 5.3.1
(Lozinskii [122]) Let Γ be a consistent set of formulae in n variables and let $\mathfrak{M}(\Gamma)$ denotes the collection of (equivalence classes of) models for Γ. The information value

of Γ *is defined by the following equation:*

$$I(\Gamma) = \log \frac{2^n}{|\mathfrak{M}(\Gamma)|} \tag{5.6}$$

Replacing equation (5.6) in base 2 we have:

$$I(\Gamma) = n - \log_2 |\mathfrak{M}(\Gamma)| \tag{5.7}$$

The intuitive justification of our definition is that the amount of information in a data set should be based on the logarithmic ratio between the number of non-equivalent interpretations and the number of equivalent models of the data. This is generally in agreement with the underlying idea of equation (5.1). If a data set allows us to exclude all interpretations except one as its model, then the data set has maximum information value. We also note that the definition applies only to data sets with finitely many variables. For sets with infinitely many variables we need to modify our definition since it is not meaningful to talk about the ratio between two infinite cardinals. For simplicity, we'll focus on sets in finitely many variables. We should mention that Lozinskii's definition of information value is very similar to the κ function defined by Gent, Prosser and Walsh [78] in their study of the *constrainedness* of search problems. Gent's κ function is intended to provide a quantitative measurement for an ensemble of search problems (e.g. **SAT** or graph colouring problems), to determine how hard or easy it is to find a solution for these problems.

In the context of inconsistent data it is natural to ask for a measurement of information analogous to our definition. However, unlike the approach of Aisbett and Gibbon in [6], we do not agree that inconsistent data provides no information at all. We equally reject the suggestion that inconsistent data provides the maximum amount of information since all *all* possibilities are eliminated. Instead we should take a more pragmatic approach here. What is and what isn't informative seems to depend, at least partly, on the goal of the agent in possession of the data. For a tax auditor, inconsistencies in a taxpayer's records are useful information for detecting possible fraud. Inconsistencies may also be useful in cases where they are deployed as directives to guide learning or as indicators for faulty components in a complex system. Worse still, by assigning null information value to all inconsistent data we may incur information loss. As we mentioned earlier, an important aspect of handling inconsistencies is the ability to compare and evaluate the relative merit of different inconsistent data sets. We need to have some quantitative criteria to determine whether one data set is more inconsistent or informative than another. Thus it is desirable to have a general theoretical framework for measuring both consistent and inconsistent information. In [123] Lozinskii provides such a framework.

Definition 5.3.2
(Lozinskii [123]) Let Γ be a set of formulae in n variables and $M(\Gamma)$ be the set of maximal consistent subsets of Γ. For each $\Delta \in M(\Gamma)$, if $\mathfrak{M}(\Delta)$ is the collection of (equivalence classes of) models of Δ then the collection of quasi-models is defined by:

$$\mathfrak{U}(\Gamma) \quad = \quad \bigcup \{\mathfrak{M}(\Delta) : \ \Delta \in M(\Gamma)\} \tag{5.8}$$

The information value of Γ is defined by the following equation:

$$I(\Gamma) \quad = \quad n - \log_2 |\mathfrak{U}(\Gamma)| \tag{5.9}$$

Again, the main idea behind definition (5.3.2) is that the information value of a set of formulae is determined by the logarithmic ratio between the number of non-equivalent interpretations and the number of *quasi-models*. Clearly definition (5.3.2) agrees with definition (5.3.1) when Γ is consistent and yields a defined value for $I(\Gamma)$ when $M(\Gamma)$ is non-empty. We note that according to the new definition the information value of a data set is monotonically increasing with respect to consistent supersets, i.e. for any consistent $\Gamma' \supseteq \Gamma$, $I(\Gamma) \leq I(\Gamma')$. For inconsistent sets however, the information value is nonmonotonic when there is an increase in inconsistencies. For instance,

Example 5.3.1
For $\Delta = \{p \vee q, \ p \vee \neg q, \ \neg p \wedge r\}$, $\Gamma = \Delta \cup \{\neg r\}$ and $\Gamma' = \Delta \cup \{s\}$

$$I(\Gamma) < I(\Delta) \qquad\qquad I(\Gamma') > I(\Delta)$$

5.4 QC Logic and Information Measure

Lozinskii's new definition is problematic in two respects. The first is that the presence of tautologies will affect the value of $I(\Gamma)$. Since we are primarily interested in the amount of *empirical* information about the world, it seems reasonable to disregard tautological statements in a data set. In a more general setting, of course, we may relativise the information value of a data set by nominating a particular set of formulae to be disregarded. This is a useful generalisation since, as we have already pointed out, the information value of a data set is at least partly dependent on the agent in possession of the data. Perhaps an agent has already independently confirmed A and thus it is not informative to be told A again. The second problem is that $I(\Gamma)$ is too sensitive to the syntax of the formulae in Γ and thus may produce counter-intuitive consequences. Indeed this is a general problem with any inconsistency tolerant mechanism based on maximal consistent subsets. The syntactic features of the formulae in the set determine how the set can be fragmented into consistent subsets. In [182], Wong considers the following example:

	$\Gamma_1 = \{p \wedge q,\ \neg p \wedge r\}$	$\Gamma_2 = \{p \wedge q \wedge r,\ \neg p \wedge q \wedge r,\ p \wedge \neg q \wedge r\}$
$\lvert M(\Gamma)\rvert$	2	3
$\lvert \mathfrak{U}(\Gamma)\rvert$	4	3
$I(\Gamma)$	1.00	1.42

Table 5.1: Information and Inconsistencies

In this example, Γ_2 is in some sense more inconsistent than Γ_1; yet we have $I(\Gamma_2) > I(\Gamma_1)$. Intuitively, the information value of a set should vary inversely to the amount of inconsistency in the set. The information value of a highly inconsistent data set should be lower than that of a set with fewer inconsistencies. A natural solution is to relativise the information value of a set using the decomposition closure defined in the previous section; that is, we let

$$I^*(\Gamma) \quad = \quad n - \log_2 \lvert \mathfrak{U}(C_D(\Gamma)) \rvert \tag{5.10}$$

Since C_D is a variable and finiteness preserving closure operator, replacing Γ with $C_D(\Gamma)$ in (5.9) has no effect on the value of n. Indeed the advantage of (5.10) over (5.9) is that it provides a more discriminating way of evaluating the information value of a data set. This gives us a more realistic appraisal of the usefulness of our data. The information value of a set no longer depends on how the formulae are syntactically presented.

Example 5.4.1
$\Gamma = \{p \vee q,\ p \vee \neg q,\ \neg p \wedge r\}$ and $\Gamma' = \{p \vee q,\ p \vee \neg q,\ \neg p,\ r\}$

Using equation (5.9) we have $I(\Gamma) \neq I(\Gamma')$. According to equation (5.10) however we have $I^*(\Gamma) = I^*(\Gamma')$. In the extreme case when $p_i \in C_D(\Gamma)$ and $\neg p_i \in C_D(\Gamma)$ for every variable p_i occurring in Γ, we have $I^*(\Gamma) = 0$ since the number of quasi-models for Γ is exactly 2^n where n is the number of distinct variables in Γ. In one sense C_D gives us a syntactic normal form for a set of formulae. Looking at our previous example, it is easy to see that I^* provides a more appropriate information value for Γ_1 and Γ_2.

	$C_D(\Gamma_1)$	$C_D(\Gamma_2)$
$\lvert M(C_D(\Gamma))\rvert$	2	4
$\lvert \mathfrak{U}(C_D(\Gamma))\rvert$	2	4
$I^*(\Gamma)$	2.00	1.00

Table 5.2: Comparison of Γ_1 and Γ_2

We note that we have not made full use of QC logic here. Indeed this is unnecessary and undesirable since the composition rule $\vee$–I allows the introduction of arbitrary new propositional variables. Clearly the introduction of new variables would interfere with the information value of a data set. In addition, we can also consider using C_D in conjunction with inference mechanisms based on maximal consistent subsets (chapters 2 and 3). The idea is similar in that we can first apply C_D to obtain a normal form for an inconsistent set and then use further inference mechanisms to extract conclusions from the set.

5.5 The Number of Q-Models

In this section, we'll address the question of how to compute the number of *quasi-models* of an inconsistent set. We can represent all possible quasi-models of Γ using the following scheme:

$$\begin{aligned} F &= \bigvee_{\mathcal{A} \in M(\Gamma)} \bigwedge_{C \in \mathcal{A}} C \\ &= (C_1^1 \wedge \ldots \wedge C_r^1) \vee \ldots \vee (C_1^{|M(\Gamma)|} \wedge \ldots \wedge C_s^{|M(\Gamma)|}) \end{aligned} \quad (5.11)$$

where each C_y^x is a reduced clause and $|M(\Gamma)|$ is the number of maximal consistent subsets of Γ. It is easy to see that the number of assignments verifying F is precisely $|\mathfrak{U}(\Gamma)|$ since an assignment v is model of F iff it is a quasi-model of Γ. Thus we can compute the number of quasi-models for Γ by counting the number of assignments verifying F. To do this however we need to observe the following:

Proposition 5.5.1

1. *If* Γ *is a set of clauses in* n *variables, then any maximal consistent subset of* Γ *has exactly* n *variables.*

2. *There are exactly* n *distinct variables in each disjunct of* F*.*

3. *The number of models for* F *is exactly the sum of the number of models for each disjunct of* F*.*

Proof:

(1) Since each $\mathcal{A} \in M(\Gamma)$ is a subset of Γ, the number of variables occurring in any $\mathcal{A} \in M(\Gamma)$ cannot be greater then n. Suppose to the contrary that some $\mathcal{A} \in M(\Gamma)$ has less than n variables. Let l_i be a literal whose variable does not occur in $\mathcal{A}$ but occurs in Γ. Then either $l_i \in \Gamma$ or l_i occurs as a disjunct of a clause $C \in \Gamma$. In both cases, we can find an assignment that satisfies $\mathcal{A} \cup \{l_i\}$ and $\mathcal{A} \cup \{C\}$, but this contradicts the maximal consistency of $\mathcal{A}$.

(2) Since each disjunct of F is a conjunction of all formulae of a maximal consistent subset of Γ, it follows from (1) that there must be n distinct variables occurring in each disjunct of F.

(3) This follows from the fact that the disjuncts of F are pairwise inconsistent and that each disjunct of F is consistent. ∎

From (3) of the above proposition, it suffices to calculate the number of models for each disjunct of F and then sum them. Now consider the k-th disjunct of F. Suppose it is of the form,

$$D^k = C_1^k \wedge \ldots \wedge C_m^k$$

We can calculate the number of assignments which verify D^k by counting the number of assignments which verify $\neg D^k$. From (2) of the above proposition, there must be n variables in D^k (respectively $\neg D^k$). So there must be 2^n distinct assignments over D^k (respectively $\neg D^k$). Where $|C_i^k|$ is the number of distinct variables in the i-th clause of D_k, the number distinct assignments which verify $\neg C_i^k$ is then given by the equation

$$|\{v|\ v \models \neg C_i^k\}| = 2^{n-|C_i^k|} \tag{5.12}$$

In general, the size of the union of a given family of sets $S_1 \cup S_2 \cup \ldots \cup S_m$ is given by the Inclusion-Exclusion formula ([40; 121]):

$$|\bigcup_{h=1}^{m} S_h| = \sum_{1\leq h\leq m} |S_h| - \sum_{1\leq h<i\leq m} |S_h \cap S_i| + \sum_{1\leq h<i<j\leq m} |S_h \cap S_i \cap S_j| - \cdots + (-1)^m |S_1 \cap \ldots \cap S_m| \quad (5.13)$$

So on the basis of (5.13), the total number of assignments which verify $\neg D^k$ is

$$|\{v| \, v \models \neg D^k\}| = \sum_{1\leq i\leq m} |\{v| \, v \models \neg C_i^k\}| - \sum_{1\leq i<j\leq m} |\{v| \, v \models \neg C_i^k \wedge \neg C_j^k\}| + \cdots + (-1)^m |\{v| \, v \models \neg C_1^k \wedge \ldots \wedge \neg C_m^k\}| \quad (5.14)$$

So the number of assignments which verify D^k is

$$\begin{aligned} |\{v| \, v \models D^k\}| &= 2^n - |\{v| \, v \models \neg D^k\}| \\ &= 2^n - \Big(\sum_{1\leq i\leq m} |\{v| \, v \models \neg C_i^k\}| - \\ &\quad \sum_{1\leq i<j\leq m} |\{v| \, v \models \neg C_i^k \wedge \neg C_j^k\}| + \cdots \\ &\quad + (-1)^m |\{v| \, v \models \neg C_1^k \wedge \ldots \wedge \neg C_m^k\}| \Big) \end{aligned} \quad (5.15)$$

The number of assignments which verify F is simply the sum of assignments which verify some disjunct of F,

$$\begin{aligned} |\{v| \, v \models F\}| &= \Big(2^n - |\{v| \, v \models \neg D^1\}|\Big) + \cdots + \Big(2^n - |\{v| \, v \models \neg D^{|M(\Gamma)|}\}|\Big) \\ &= |M(\Gamma)| \cdot 2^n - \sum_{1\leq i\leq |M(\Gamma)|} |\{v| \, v \models \neg D^i\}| \end{aligned} \quad (5.16)$$

We note that in equation (5.14), if C_i^k and C_j^k are clauses which contain complementary literals p and $\neg p$, then $\neg C_i^k \wedge \neg C_j^k$ is unsatisfiable and thus $|\{v| \, v \models \neg C_i^k \wedge \neg C_j^k\}| = 0$. Moreover any conjunctive extension of $\neg C_i^k \wedge \neg C_j^k$ would clearly also be unsatisfiable. Thus one or more terms of the **RHS** of (5.14) may turn out to be zero. On the other hand if $C_1^k \ldots C_m^k$ are pairwise free from complementary literals, then $\neg C_1^k \wedge \ldots \wedge \neg C_m^k$

is clearly satisfiable. As usual if we represent a clause C as a set of literals, the number of models for $\neg C_1^k \wedge \ldots \wedge \neg C_m^k$ can be given by

$$|\{v|\ v \models (\neg C_1^k \wedge \ldots \wedge \neg C_m^k)\}| \quad = \quad 2^{n-(|C_1^k \cup \ldots \cup C_m^k|)} \tag{5.17}$$

These observations provide the basis of an algorithm developed by Lozinskii in [121] for computing the number of models of a CNF formula.

Algorithm 5.5.1 Lozinskii's algorithm

Require: input $CNF(E)$
Ensure: output $\mu CNF(E) = |\{v|\ v \models CNF(E)\}|$

1: $s := 1;$
2: $G_1 := \{\{C\}|C \in CNF(E)\};$
3: $t_1 := \sum_{\{C_i \in G_1\}} 2^{n-|C_i|};$
4: $acc := t_1;$
5: **while** $G_s \neq \emptyset$ **do**
6:     $s := s + 1;$
7:     $G_s := \emptyset;$
8:     $t_s := 0;$
9:     **for all** $g_{s-1} \in G_{s-1}$ **do**
10:         **for all** $C \in CNF(E)$ and $C \notin g_{s-1}$ **do**
11:             $g_s := g_{s-1} \cup \{C\};$
12:             **if** g_s is pairwise free from complementary literals **then**
13:                 $G_s := G_s \cup \{g_s\};$
14:                 $t_s := t_s + 2^{n-|\cup g_s|};$
15:             **end if**
16:         **end for**
17:         $acc := acc + (-1)^{s-1} t_s;$
18:     **end for**
19: **end while**
20: $\mu CNF(E) := 2^n - acc$

The basic idea of algorithm (5.5.1) is to incrementally sum the terms of equation (5.14) from left to right. At the end of line (4), algorithm (5.5.1) computes the first term of equation (5.14), giving the value for $\sum_{1 \leq i \leq m} |\{v|\ v \models \neg C_i^k\}|$. The outer loop at line (5) is then executed (at most m times). Each run of the loop adds the value of the next term of the equation. The first execution of the outer loop computes the value of the second term $\sum_{1 \leq i < j \leq m} |\{v|\ v \models \neg C_i^k \wedge \neg C_j^k\}|$. If this value is non-zero, then acc is updated to provide the sum of the first two terms of (5.14). The loop is executed repeatedly until either control reaches the last term of (5.14) or when every possible way to extend every g_{s-1} at line (11) results in complementary literals appearing in some pair of clauses in the extension. At the first inner loop at line (9), G_{s-1} is the set of all subsets of $CNF(E)$ that are of size $(s-1)$ i.e. each $g_{s-1} \in G_{s-1}$ contains exactly $s-1$ clauses of $CNF(E)$. At line (10) the second inner loop extends g_{s-1} with a single clause from $CNF(E)$ and the resulting extension is tested for complementary literals. If the extension is pairwise free from complementary literals, then G_s is extended to include g_s and t_s is recalculated to reflect the change. Once all possible ways to extend a given g_{s-1} have been examined, control exits from the inner most loop and acc is updated to reflect the new value. Control then returns to the top of the loop at line (9) and selects another member of G_{s-1} and continues the process of the inner-most loop.

For a set of clauses Γ with a known number of maximal consistent subsets k, counting quasi-models of Γ can be done by calling Lozinskii's algorithm k times. But other propositional model counting algorithms can also be deployed to achieve the task. Indeed, experimental and theoretical results show that Lozinskii's algorithm is superseded by the Counting Davis-Putnum algorithm (CDP) of Birnbaum and Lozinskii [40] and by the Decomposing Davis-Putnum algorithm (DDP) of Bayardo and Pehoushek [18]. But we should point out that the general problem of counting propositional models (**#SAT**) is **#P**-complete. So according to the present state-of-the-art, the problem is computationally intractable in the worst case. Counting quasi-models is no easier.

5.6 Application

In previous works Hunter and Nuseibeh [93; 94] have illustrated the usefulness of QC logic in the analysis of inconsistent specifications in software engineering. Hunter and Nuseibeh have pointed out that inconsistent specifications are often unavoidable during software development. They argued persuasively that during the software development cycle it is often more important to manage inconsistencies intelligently, i.e., we need to analyse and to keep track of inconsistencies rather then resolving them immediately. In the same spirit we advocate using QC logic and the definition given by equation (6) as a basis to analyse over-constrained problems.

5.6.1 Constraint Satisfaction Problems

A constraint satisfaction problem (**CSP**) involves,

1. a set of variables, $X_1, \ldots, X_n$

2. associated with each variable, X_i, is a domain, D_i of values

3. a set of constraints, $C_1, \ldots, C_m$, each is defined on subset of variables over a subset of the Cartesian product of the associated domains, i.e. $C_i(X_{i_1}, \ldots, X_{i_k}) \subseteq (D_{i_1} \times \ldots \times D_{i_k})$

A *solution* to a **CSP** is simply an assignment of values to variables such that all constraints are satisfied. A **CSP** is a Finite Constraint Satisfaction Problem (**FCSP**) if its constraint domains are finite. Many real world problems such as optimization problems or job scheduling problems can be viewed as **CSP**s.

As is well known, there is a close relationship between **FCSP**s and logic (see [39; 125]). Any **FCSP** can be stated as an equivalent logic problem in a variety of settings. In the model checking approach for instance, a **FCSP** is taken to have a solution iff a certain propositional theory Γ is satisfiable. In fact, the solutions are just the set of models of Γ. In this scheme, the theory Γ is constructed as a set of propositional formulae in CNF such that

1. Each possible combination of values for the variables is represented by a set of propositional variables, $p^{x_1}_{d_1}, \ldots, p^{x_m}_{d_n}, \ldots$, where intuitively, $p^{x_i}_{d_j}$ is the proposition which says that the variable x_i is instantiated by the value d_j. For instance, the sentence $(p^{x_i}_{d_j} \vee p^{x_i}_{d_k})$ says that the variable x_i is instantiated by at least one of values d_j and d_k.

2. A constraint is stated *negatively* in terms of values that are forbidden, e.g. the sentence $\neg p^{x_i}_{d_j}$ says that x_i is never instantiated by value d_i, the sentence $\neg p^{x_i}_{d_k} \vee \neg p^{x_j}_{d_k}$ says that x_i and x_j are never instantiated by the same value d_k. The set of all constraints is represented by a set of propositional formulae in the variables, $p^{x_1}_{d_1}, \ldots, p^{x_m}_{d_n}, \ldots$.

5.6.2 Over-constrained Problems

As it is with many real world problems, a **CSP** can be without a solution. A solutionless **CSP** is an over-constrained problem (**OCP**) – every assignment of values to variables fails to satisfy at least one constraint. Consider for instance,

Example 5.6.1
Let X, Y, Z be variables whose domain is $\{1, 2, 3\}$. Let the constraints be: $X < Y$, $Y < Z$ and $Z < X$.

Clearly, this is an **OCP** since no natural numbers can satisfy all three constraints. This example illustrates that there are two main factors which contribute to a problem being over-constrained – and thus provides two different approaches to resolving **OCP**s. The first is the domain of possible values and the second is the constraints themselves. In our example if we were to add a value w to the domain such that for some m and n, $m < n$, $n < w$, and $w < m$, then all constraints would be satisfied (w need not be a natural number), in which case we no longer have an over-constrained problem. Alternatively, we may accept a certain *partial* assignment that satisfies some but not all of the constraints as a solution. Typically, we may accept those assignments that satisfy a maximal number of constraints or variables. Given that any **FCSP** is equivalent to a model checking problem in propositional logic, the second approach to solving a finite **OCP** is equivalent to finding models for a certain subset of an inconsistent set of propositions.

Regardless of how we may resolve an **OCP**, it is sometime desirable to analyse the problem first before any further action is taken. In this respect, it is clear that QC logic is well suited to the task. According to our previous scheme, we can encode a finite **OCP** as a propositional theory Γ; Γ must be unsatisfiable and thus inconsistent. We can then apply QC logic to analyse the information value of Γ. In particular in an **OCP** not all variables may be involved in an inconsistency (i.e. being overly constrained). Thus it is desirable to identify those variables that are involved in an inconsistency. The strategy, as before, is to take the decomposition closure of Γ and then measure the value $I^*(\Gamma)$. In a highly over-constrained problem we should expect to see a lower value for $I^*(\Gamma)$ and vice versa. This gives us a relative measurement of the constrainedness of **OCP**s.

5.7 Conclusion

In this chapter we have argued that there are general advantages in developing practical reasoning systems that can tolerate inconsistencies. In this respect we have considered a paraconsistent logic that can avoid drawing trivial conclusions in the presence of inconsistencies. But more importantly we advocate the use of paraconsistent logic in assisting us in analysing inconsistent data. In this light, the role of logic goes beyond capturing valid forms of inference. Logic can be seen as a tool for analysis.

Chapter 6

Modalized Inconsistencies

6.1 Introduction

In the standard Kripkean binary relational semantics, the truth condition for modal formulae is defined by

$$\models^{\mathfrak{M}}_{x} \Box A \;\Leftrightarrow\; \forall y,\; \mathcal{R}xy \;\Rightarrow \models^{\mathfrak{M}}_{y} A$$

where the notions of frame, model, and satisfaction are defined in the usual way (see [48; 88; 89; 90]).

The minimal modal logic determined by the Kripkean binary relational frame is the logic K, most economically axiomatised by adding to propositional logic PL the single rule called Scott Rule:

$$[SR]\ \frac{\Gamma \vdash A}{\Box[\Gamma] \vdash \Box A}$$

where $\Box[\Gamma] = \{\Box B :\ B \in \Gamma\}$. Alternatively, K can be axiomatised by adding to PL the axiom schemata

$$[K]\ \Box A \wedge \Box B \rightarrow \Box(A \wedge B)$$

$$[N]\ \Box\top$$

and the rule of monotonicity

$$[RM]\ \frac{A \rightarrow B}{\Box A \rightarrow \Box B}$$

In [164; 165; 100; 101], Jennings and Schotch generalise the Kripkean binary relational semantics for modal logics by replacing the Kripkean binary relation with an $(n+1)$-ary relation for each $n > 1$. The truth condition for modal formulae is then redefined as

$$\models^{\mathfrak{M}}_{x} \Box A \;\Leftrightarrow\; \forall y_1, \ldots, y_n,\; \mathcal{R}xy_1, \ldots, y_n \;\Rightarrow\; \exists i_{(1\leq i\leq n)}:\; \models^{\mathfrak{M}}_{y_i} A$$

Just as K is the minimal logic associated with the class of Kripkean binary relational frames, associated with each class of $(n+1)$-ary frames is a minimal weak aggregative modal logic K_n. For each $n > 1$, the logic K_n is obtained by replacing [K] with the weaker aggregation schema

$$[K_n]\ \Box A_1 \wedge \ldots \wedge \Box A_{n+1} \rightarrow \Box(\bigvee_{1 \leq i < j \leq n+1} A_i \wedge A_j)$$

We have, in fact, a descending sequence of modal logics ordered by inclusion:

$$K \supseteq K_2 \supseteq K_3 \supseteq \ldots$$

The limit of such a sequence turns out to be $\bigcap_{n<\omega} K_n = N$, which is finitely axiomatisable by PL, [N] and [RM] alone. As is well known, there are modal logics even weaker than N, for instance E which has the single rule

$$[RE]\ \frac{A \leftrightarrow B}{\Box A \leftrightarrow \Box B}$$

In [11], Apostoli and Brown deploy a general strategy to show that for each $n > 1$, K_n is determined by its respective class of $(n+1)$-ary relational frames, and thereby verify a claim made in [164]. Their strategy hinges on showing that K_n can be alternatively axiomatised by the single rule

$$[RT_n]\ \frac{\Gamma \vdash_n A}{\Box[\Gamma] \vdash \Box A}$$

where $\Gamma \vdash_n A$ just iff every n-partition of Γ contains a cell which classically entails A.

The inclusion ordering of various weak modal logics is summarised in figure (6.1). The logic M adopts the rule [RM]; C adopts [RM] and [K]; and N^* adopts the rule of necessitation: [RN] $\vdash A \Longrightarrow \vdash \Box A$. N^{**} adopts a weakened version of [RN] without the restriction that A is a theorem, i.e. necessitation applies to any formula. The logic N^{**} turns out to have some surprisingly deep connections with default logic (see [70]).

The initial motivation to study K_n logics comes partly from the interest in finding a suitable medium to express deontic, doxastic and epistemic dilemmas ([101; 166]). In the presence of the strong aggregation principle [K], no distinction can be made between having several incompatible obligations (moral, legal etc.) and having an obligation to inconsistency. In doxastic and epistemic contexts where we may take '$\Box A$' to mean 'It is warranted that A' or 'It is justified that A', it is equally unreasonable

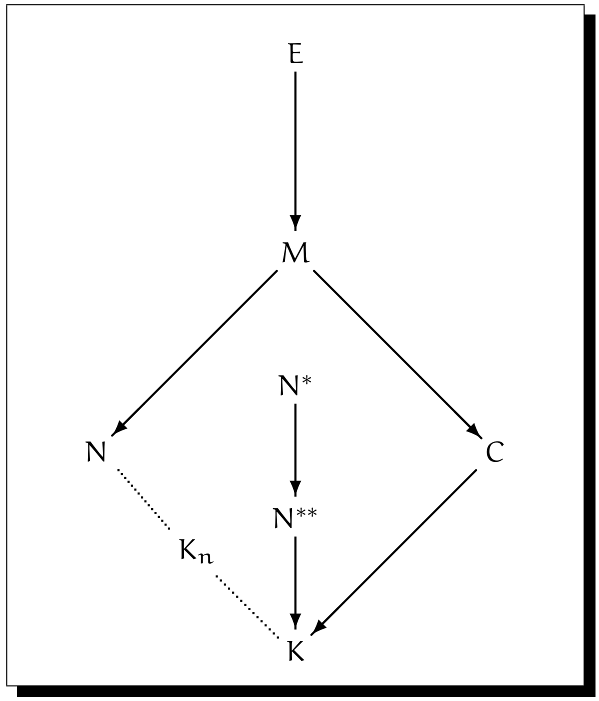

Figure 6.1: Inclusion ordering of weak modal logics.

to suppose that having incompatible but individually warranted claims amounts to an inconsistency being warranted. In each of these cases, the problem is the failure to observe the distinction between

$$[\mathrm{I}]\ \Box A \wedge \Box\neg A \quad \text{and} \quad [\mathrm{I}^*]\ \Box(A \wedge \neg A)$$

In the standard Kripkean binary relational semantics, points at which [I] is true are exactly the same points where [I*] is true - namely points from which other points are not accessible. The collapse of the distinction between [I] and [I*] leads immediately to the collapse of the further distinction between the consistency principles:

$$[\mathrm{D}]\ \Box A \to \neg\Box\neg A \quad \text{and} \quad [\mathrm{Con}]\ \neg\Box\bot$$

Arguably, in the deontic or epistemic reading of '□', [Con] is a plausible principle which requires that no inconsistency be obligatory or warranted. [D] however makes the stronger demand that incompatible obligations or claims are ruled out at the outset. Worst still, in the presence of [K] and [RM], [I] implies $\Box B$ for any B. [1] Thus in any extension of the logic K, having several incompatible obligations or having

[1]Since $A \wedge \neg A \to B$ is a tautology, we have $\Box(A \wedge \neg A) \to \Box B$, and by [K] we get $\Box A \wedge \Box\neg A \to \Box B$.

incompatible but individually warranted claims tantamounts to everything is obligatory or warranted. Of course, we may just bite the bullet and insist that dilemma, moral and the like never arise in deontically and doxastically ideal worlds. But from a design stance, if we are to make our robots reason more like us and less like God, unrestricted aggregation of deontic or epistemic modalities is not always desirable. Our robots, like us, live in a world in which dilemmas and conflicts lie abound. Thus it seems desirable to endow our robots with some capacity to reason with incompatible information or obligations.

So much for philosophical motivations. On a technical level, from both a proof theoretic and a semantic standpoint, the generalisation from binary to $(n+1)$-ary relational frames is an interesting strategy with which to study modal logics weaker than K while still remaining within some sort of relational semantics. But the deployment of multi-ary relational semantics also stages another strategy to generalise modal logics - in particular, from logics with unary modal operators to logics with multi-ary modal operators. [2]

Several examples of such generalisation are readily available. Routley and Meyer [129; 157], Gabbay [72], Johnston [102], Goldblatt [79], and Bell [19] have all independently introduced logics with multi-ary modal operators and shown their completeness with respect to several multi-ary relational frames. If '$\Box$' is now taken to be an n-ary modal operator, and

$$C = \Box(B_1, \dots, B_{i-1}, B, B_{i+1}, \dots, B_n)$$

then we may write '$C_B[A]$' to denote, the formula obtained by replacing B with A in C, i. e.

$$C_B[A] = \Box(B_1, \dots, B_{i-1}, A, B_{i+1}, \dots, B_n)$$

Having this notation available at hand, in [79] Goldblatt shows that the logic K^n, [3] axiomatised by adding to PL the axiom schemata

$$[K^n]\ C_B[A] \wedge C_B[D] \rightarrow C_B[A \wedge D]$$

[2] Of course this is not surprising since the algebraic foundations of relational semantics were articulated by Jónsson and Tarski in their seminal papers [103; 104]. According to Copland [49], Kripke had remarked that Jónsson and Tarski's paper [103] was the 'most surprising anticipation' of his own work. Routley and Meyer [157] also deserve credit for generalising binary relational semantics for modal logic to ternary relational semantics for relevant implications.

[3] Although Goldblatt's axiomatisation differs from ours, his logic is equivalent to ours.

$$[N^n]\ C_B[\top]$$

and the rule of inference

$$[RM^n]\ \frac{A_i \to B_i\ (1 \le i \le n)}{\Box(A_1, \ldots, A_n) \to \Box(B_1, \ldots, B_n)}$$

is determined by the class of $(n+1)$-ary relational frames whose truth condition for modal formulae is defined by

$$\models_x^{\mathfrak{M}} \Box(A_1, \ldots, A_n) \Leftrightarrow \forall y_1, \ldots, y_n,\ \mathcal{R}xy_1, \ldots, y_n \Rightarrow\ \exists j_{(1\le j\le n)}:\ \models_{y_j}^{\mathfrak{M}} A_j$$

We note that Goldblatt's truth condition is almost but not quite the same as Jennings-Schotch's. In Jennings-Schotch's semantics, $\Box A$ is true at a point just in case A is true somewhere in every related n-tuple of points. In Goldblatt's semantics, $\Box(A_1, \ldots, A_n)$ is true at a point just in case each related n-tuple of points has some j, $1 \le j \le n$, where A_j is true in the j-th coordinate of the n-tuple. In each case, the Kripkean truth condition is uniquely recovered when we set $n = 1$.

To no one's surprise, just as K can be axiomatised by [SR], Goldblatt's K^n can be axiomatised by the rule

$$[GR]\ \frac{\Gamma \vdash A}{C_B[\Gamma] \vdash C_B[A]}$$

where $C_B[\Gamma] = \{C_B[D] : D \in \Gamma\}$. There is clearly a symmetry between unary modal logics and n-ary modal logics. $[K^n]$, $[N^n]$, and $[RM^n]$ are just n-ary counterparts of the familiar [K], [N], and [RM]. Similarly, the rule [RE] can be restated as an n-ary counterpart:

$$[RE^n]\ \frac{A_i \leftrightarrow B_i\ (1 \le i \le n)}{\Box(A_1, \ldots, A_n) \leftrightarrow \Box(B_1, \ldots, B_n)}$$

Our main purpose here is to study a generalisation of Jennings-Schotch's logics. In particular, we'll show that just as normal unary modal logic K can be weakened to K_n by progressively relaxing the aggregation principle, multi-ary modal logic K^n can also be weakened to the logics K_n^m. Such logics can in fact be axiomatised by a rule analogous to $[RT_n]$. We'll generalise the Apostoli-Brown strategy to show that K_n^m is determined by a class of $m+n$-ary relational frames. In section (6.2), we'll introduce our logics and their semantics. In section (6.3), we'll introduce a species of paraconsistent consequence relations, called n-forcing, based on our logics and prove the compactness property for such consequence relations. In section (6.4), we'll present Apostoli-Brown's axiomatisation of n-forcing and its completeness. In section (6.5), we'll give completeness a proof for our logics. We'll conclude with a conjecture.

6.2 Logical Preliminaries

6.2.1 Syntax

A set of formulae, Φ, is constructed in the usual way from a set of propositional atoms, $At = \{p_1, p_2, \ldots\}$, and a set of connectives, $\neg, \wedge, \vee, \Box$ where $\Box$ is an m-ary connective. As usual we'll omit outermost parenthesis. We'll use $A \to B$ and $A \leftrightarrow B$, $\bot$ and $\top$ as shorthand for $\neg A \vee B$, $(\neg A \vee B) \wedge (\neg B \vee A)$, $A \wedge \neg A$, $A \vee \neg A$, respectively. For each $n \geq 1$ and each $m \geq 1$, the modal logic K_n^m ($\Lambda \subseteq \Phi$) is the least set satisfying the following conditions:

- $PL \subseteq \Lambda$ and Λ is closed under the rules of PL
- where $C = \Box(C_1, \ldots, C_{m-1}, B)$, $[K_n^m] \in \Lambda$, i. e.

$$C_B[A_1] \wedge \ldots \wedge C_B[A_{n+1}] \to C_B[\bigvee_{1 \leq i < j \leq n+1} A_i \wedge Aj] \in \Lambda$$

- $[N_n^m] \in \Lambda$ i. e.

$$C_B[\top] \in \Lambda$$

- Λ is closed under $[RM_n^m]$, i. e. for every i, $1 \leq i \leq m-1$,

$$A_i \leftrightarrow B_i \in \Lambda \text{ and } A_m \to B_m \in \Lambda \Longrightarrow \Box(A_1, \ldots, A_m) \to \Box(B_1, \ldots, B_m) \in \Lambda$$

If A is a theorem of Λ, we write, $\vdash_\Lambda A$. And for any $\Gamma \subseteq \Phi$, $\Gamma \vdash_\Lambda A$ iff there is an $n \in \mathbb{N}$ such that $B_1, \ldots, B_n \in \Gamma$ and $\vdash_\Lambda B_1 \wedge \ldots \wedge B_n \to A$. A set Γ is Λ-inconsistent iff $\Gamma \vdash_\Lambda \bot$. Where the context is clear, we'll use 'consistent' instead of Λ-consistent.

6.2.2 Models

A model $\mathfrak{M} = \langle \mathcal{U}, \mathcal{R}, \mathcal{V} \rangle$ where $\mathcal{U} \neq \emptyset$, $\mathcal{R} \subseteq (\wp(\mathcal{U}))^{m-1} \times \mathcal{U}^{n+1}$ and $\mathcal{V} : At \longrightarrow \wp(\mathcal{U})$. The satisfaction relation $\models$ is defined inductively by

- $\models_x^{\mathfrak{M}} p_i \;\Leftrightarrow\; x \in \mathcal{V}(p_i)$
- $\models_x^{\mathfrak{M}} \neg A \;\Leftrightarrow \not\models_x^{\mathfrak{M}} A$
- $\models_x^{\mathfrak{M}} (A \vee B) \;\Leftrightarrow \models_x^{\mathfrak{M}} A$ or $\models_x^{\mathfrak{M}} B$

- $\models_x^{\mathfrak{M}} (A \wedge B) \Leftrightarrow \models_x^{\mathfrak{M}} A$ and $\models_x^{\mathfrak{M}} B$

- $\models_x^{\mathfrak{M}} \Box(A_1, \ldots, A_{m-1}, B) \;\Leftrightarrow\; \forall y_1, \ldots, y_n \in \mathcal{U},$

$$\mathcal{R}\|A_1\|^{\mathfrak{M}}, \ldots, \|A_{m-1}\|^{\mathfrak{M}}, xy_1, \ldots, y_n \Rightarrow \exists i_{(1\leq i\leq n)}\colon \models_{y_i}^{\mathfrak{M}} B$$

where for any $C \in \Phi$, $\|C\|^{\mathfrak{M}} = \{x \in \mathcal{U} : \models_x^{\mathfrak{M}} C\}$. We note that the Jennings-Schotch's semantics is recovered when $m = 1$, one of Gabbay's semantics as defined in [72] (p. 180) is recovered when $n = 1$. And when $m = n = 1$, then we have the standard Kripkean semantics.

Theorem 6.2.1
K_n^m *is sound with respect to our models.*

The soundness proof is standard. We leave it to the reader to verify.

6.3 $\mathfrak{n}$-Forcing and Coherence Level

Before we tackle the completeness problem of K_n^m, we'll need some additional definitions and lemmata. In this section, we'll introduce the n-*forcing* inference relation - a species of paraconsistent inference relation, and the notion of coherence level of a set. From now on, we'll assume that m is fixed. Relative to Λ, the notion of n-forcing and coherence level are defined as follows:

Definition 6.3.1
A non-empty Λ-inconsistent set Γ n-forces A, $\Gamma \vdash_n A$ iff for every n-partition, π, of Γ, there is a cell, $\mathcal{C} \in \pi$, such that $\mathcal{C} \vdash_\Lambda A$. If $\Gamma = \emptyset$ or is Λ-consistent, then $\Gamma \vdash_n A$ iff $\Gamma \vdash_\Lambda A$.

The collection of all n-partitions of Γ will be denoted by $\prod_n(\Gamma)$. We'll say that a partition of Γ is a Λ-consistent partition iff each cell of the partition is Λ-consistent.

Definition 6.3.2
The coherence level of a set Γ, $\ell : \wp(\Phi) \longrightarrow \mathbb{N} \cup \{\infty\}$ *is defined as follows:*

$$\ell(\Gamma) = \begin{cases} 1 & \text{if } \Gamma = \emptyset \text{ or } \Gamma \nvdash_\Lambda \bot \\ \text{the cardinality of the least } \Lambda\text{-consistent partition of } \Gamma \text{ up to and including } \omega & \text{if such partition exists} \\ \infty & \text{otherwise} \end{cases}$$

Given the usual notion of a maximal consistent set, we can state more explicitly the relationship between n-forcing and coherence level of a set and Λ-maximal consistent sets. We'll use $[\wedge C]\Gamma$ to denote $\{B \wedge C : B \in \Gamma\}$, if $\Gamma = \emptyset$, then we let $[\wedge C]\Gamma = \{C\}$. Where $\Sigma = [\wedge C]\Gamma$, we'll let $[\wedge C]^*\Sigma = \{B : B \wedge C \in \Sigma\}$.

Proposition 6.3.1
The following statements are equivalent

1. $\Gamma \vdash_n A$
2. $\ell([\wedge\neg A]\Gamma) > n$
3. *For any maximal* Λ*-consistent sets* $x_1, \ldots, x_n$ *such that* $\Gamma \subseteq \bigcup\{x_1, \ldots, x_n\}$, $A \in \bigcup\{x_1, \ldots, x_n\}$.

Proof:
(1)⇒(2): Assume that $\Gamma \vdash_n A$. If $\ell(\Gamma) > n$, clearly $\ell([\wedge\neg A]\Gamma) > n$. So assume that $\ell(\Gamma) \leq n$. Towards a contradiction, assume that $\ell([\wedge\neg A]\Gamma) \leq n$. Then there must be a consistent n-partition of $[\wedge\neg A]\Gamma$. Let $\pi = \{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ be such a consistent n-partition. Then $\pi^* = \{[\wedge\neg A]^*\mathcal{C}_1, \ldots, [\wedge\neg A]^*\mathcal{C}_n\}$ is a consistent n-partition of Γ. But $\Gamma \vdash_n A$, so $\exists i_{(1\leq i\leq n)}$: $[\wedge\neg A]^*\mathcal{C}_i \vdash_\Lambda A$. This contradicts our assumption that every $\mathcal{C}_i$ is Λ-consistent. So $\ell([\wedge\neg A]\Gamma) > n$ as required.

(2)⇒(3): Assume that $\ell([\wedge\neg A]\Gamma) > n$. Let $x_1, \ldots, x_n$ be any maximal Λ-consistent sets such that $\Gamma \subseteq \bigcup\{x_1, \ldots, x_n\}$. Clearly there must be a consistent n-partition of Γ such that each cell, $\mathcal{C}_i$, is a subset of x_i. Let $\pi = \{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ be such a consistent n-partition. Then $\pi' = \{[\wedge\neg A]\mathcal{C}_1, \ldots, [\wedge\neg A]\mathcal{C}_n\}$ is an n-partition of $[\wedge\neg A]\Gamma$. But by our initial assumption $\ell([\wedge\neg A]\Gamma) > n$, so $\exists i_{(1\leq i\leq n)} : \mathcal{C}_i \vdash_\Lambda A$. Hence, $\exists i_{(1\leq i\leq n)} : x_i \vdash_\Lambda A$.

By deductive closure of maximal consistent sets, $A \in x_i$. Hence $A \in \bigcup\{x_1, \ldots, x_n\}$ as required.

(3)⇒(1): Assume that $\Gamma \not\vdash_n A$. Clearly, if $\ell(\Gamma) > n$, then $\Gamma \vdash_n A$. So $\ell(\Gamma) \leq n$. Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ be a consistent n-partition such that $\forall i_{(1 \leq i \leq n)} : \mathcal{C}_i \not\vdash_\Lambda A$. Such a partition clearly exists, otherwise $\Gamma \vdash_n A$. Then $\forall i_{(1 \leq i \leq n)}$, $\mathcal{C}_i \cup \{\neg A\}$ is Λ-consistent. We extend each $\mathcal{C}_i \cup \{\neg A\}$ to its maximal Λ-consistent extension. Hence there exist n maximal Λ-consistent sets $x_1, \ldots, x_n$ such that $\Gamma \subseteq \bigcup\{x_1, \ldots, x_n\}$, but $A \notin \bigcup\{x_1, \ldots, x_n\}$. ∎

Calling $\vdash_n$ a consequence relation is well suited since it satisfies the usual properties of reflexivity, monotonicity, and transitivity (we'll leave it to the reader to verify this). And as we show later, it is also finitary. However, stepping back from the particular logic Λ and looking at things a bit more abstractly, proposition (6.3.1) underscores the fact that $\vdash_n$ generalises the classical consequence relation $\vdash$. The classical counterpart to proposition (6.3.1) is the familiar equivalence between (1) $\Gamma \vdash A$, (2) $\Gamma \cup \{\neg A\}$ is inconsistent, and (3) $A \in x$ for any maximal consistent extension x of Γ.

Another related generalisation at work is the notion of coherence level. In this framework, classically consistent sets are simply level 1 sets whereas all classically inconsistent sets are level n sets, where $n \geq 2$. Thus a classically consistent theory, in the sense of a consistent deductively closed set, is simply a level 1-theory closed under classical $\vdash$. It is not difficult to see that just as the closure of a level 1 set under classical $\vdash$ yields a level 1-theory, closure of a level n set under $\vdash_n$ yields a level n-theory. Thus we may say that $\vdash_n$ is a level preserving relation for any set with level $\leq n$. Now a theory is said to be *trivial* iff every formula is a deductive consequence of the theory. As is well known, closing a level $n \geq 2$ set under classical $\vdash$ yields a trivial theory. Thus as some put it colourfully, classical $\vdash$ is inferentially explosive with respect to inconsistent sets. Given these observations, $\vdash_n$ provides a possible strategy for studying inconsistent but non-trivial theories as well as paraconsistent formal systems in which not every B is a deductive consequence of $\{A, \neg A\}$. More interestingly, from an information processing viewpoint n-forcing provides a plausible inferential strategy to extract information from multiple sources, where n corresponds to the number of information channels or sources. More detailed discussions of n-forcing and inconsistency-tolerant reasoning can be found in [99] and [167].

One of the crucial steps in Apostoli-Brown's proof of the completeness of K_n is to establish the compactness of $\vdash_n$. But proposition (6.3.1) now makes it clear that the compactness of $\vdash_n$ is an immediate corollary of the compactness of the coherence levels of sets. We'll state the problem of compactness of coherence level in terms of *trace*, a kind of generalised filter base, as presented by Jennings and Schotch in [167]. The notion of trace, we may add, is in fact equivalent to the notion of non-colourability of hypergraphs. [4]

Definition 6.3.3
Let Σ be a collection of finite subsets of a non-empty set Γ. Then Σ is an n-trace over Γ iff for every n-partition, π, of Γ, there is a cell $\mathcal{C} \in \pi$ such that some element of Σ is a subset of $\mathcal{C}$.

Lemma 6.3.1
If Σ and Γ are non-empty finite sets and Σ is an n-trace over Γ, then Σ is an m-trace over Γ, for $m < n$.

Proof:
We assume that Σ and Γ are non-empty, finite and Σ is an n-trace over Γ. Let $\Sigma = \{\mathcal{A}_i : 1 \leq i \leq k\}$. Suppose for some $m < n$, Σ is not an m-trace over Γ. Then there is an m-partition $\pi = \{\mathcal{C}_j : 1 \leq j \leq m\}$ of Γ such that $\mathcal{A}_i \not\subseteq \mathcal{C}_j$ for all j. Let

$$\pi^* = \{\mathcal{B}_l : 1 \leq l \leq n, \ \emptyset \neq \mathcal{B}_l \subseteq \mathcal{C}_j, \text{ for some } j\}$$

Then π^* is an n-partition such that $\mathcal{A}_i \not\subseteq \mathcal{B}_l$ for each $\mathcal{A}_i$ and each $\mathcal{B}_l$. But this contradicts the hypothesis that Σ is an n-trace over Γ. Hence, for each $m < n$, Σ is an m-trace over Γ. ∎

Lemma 6.3.2
Let $\Gamma \neq \emptyset$ and Σ be an n-trace over Γ. Then $\exists \Sigma_0 \subseteq_{fin} \Sigma$ such that Σ_0 is an n-trace over Γ.

Proof:
Our strategy is to prove the contrapositive, i.e. if every finite subset of Σ is not an n-trace over Γ, then Σ is not an n-trace over Γ. We proceed to construct a first order theory T such that every finite subset of Σ is not an n-trace of Γ iff every finite subset of T has a model. So by the compactness of first order logic, if every finite subset of Σ

[4] A hypergraph, G, is a pair, $(V(G), E(G))$ where $V(G)$ is a set of *vertices*, and $E(G)$ is a collection of *edges*, i.e. a collection of *finite* subsets of $V(G)$. An n-colouring of a hypergraph G is an n-partition of $V(G)$. A hypergraph G is n-colourable iff some n-colouring c of G is such that no edge of G is monochromatic under c, i.e. no edge is a subset of any cell of the partition c. The compactness theorem for hypergraphs states that a hypergraph G is n-colourable iff every finite sub-hypergraph of G is n-colourable. In proving the compactness of trace, we thereby also establish the compactness of hypergraphs.

is not an n-trace of Γ, then T has a model and hence by the construction of T, Σ is not an n-trace over Γ. Let

$$\Sigma = \{\mathcal{A}_i : i \in I\}$$

For each $i \in I$, let

$$\mathcal{A}_i = \{a^i_1, \ldots, a^i_{k_i}\}$$

Make the assumption that every finite subset of Σ is not an n-trace over Γ, i. e. for each finite subset Σ' of Σ there is an n-partition of Γ such that no cell in the partition contains any element of Σ'. Towards the construction of our first order theory T, we extend the first order language with

- n many predicate symbols, $P_1, \ldots, P_n$, each representing a cell in the n-partition
- for each $i \in I$, introduce constant symbols, $c^i_1, \ldots, c^i_{k_i}$ each naming the corresponding element in $\mathcal{A}_i$.

Let

$$\Theta = \forall x(\bigvee_{1\leq h\leq n} P_h x)$$

For each $i \in I$, let

$$\Omega_i = \Big\{\Big(\bigwedge_{1\leq h\leq n}(\bigvee_{1\leq j\leq k_i} \neg P_h c^i_j)\Big) \wedge \Theta\Big\}$$

Let

$$\Omega = \bigcup_{i\in I}\Omega_i$$

We obtain the first order theory T by adding all elements of Ω as proper axioms to standard first order logic. But our assumption is that every finite subset of Σ is not an n-trace over Γ, so every finite subset of Ω has a model. By first order compactness, T has a model and thus Ω has a model. Hence there must be an n-partition of Γ such that no cell of the partition contains any element of Σ, i.e. Σ is not an n-trace over Γ. ∎

Lemma 6.3.3
Let Σ be finite, $\Gamma \neq \emptyset$. If Σ is an n-trace over Γ, then $\exists\Gamma_0 \subseteq_{fin} \Gamma$ such that Σ is an n-trace over Γ_0.

Proof:
Assume that Σ is finite and $\Gamma \neq \emptyset$. We'll show that if Σ is not an n-trace over Γ', $\forall\Gamma' \subseteq_{fin} \Gamma$, then Σ is not an n-trace over Γ. Our strategy is similar to the proof of

lemma (6.3.2). We proceed to construct a first order theory T such that Σ is not an n-trace over Γ', $\forall\Gamma' \subseteq_{fin} \Gamma$, iff every finite subset of T has a model. So by the compactness of first order logic, if Σ is not an n-trace over Γ', $\forall\Gamma' \subseteq_{fin} \Gamma$, then T has a model and hence by the construction of T, Σ is not an n-trace over Γ. Let I be an index set. For each $i \in I$, let $\Gamma_i \subseteq_{fin} \Gamma$. Let $\Sigma = \{\mathcal{A}_1, \ldots, \mathcal{A}_m\}$. For $1 \leq j \leq m$, let $\mathcal{A}_j = \{a_1^j, \ldots, a_{k_j}^j\}$. Make the assumption that Σ is not an n-trace over any finite subset of Γ, i. e. for each $i \in I$, there is an n-partition of Γ_i such that no cell in the partition includes a_j, $1 \leq j \leq m$. Towards the construction of our first order theory T, extend the first order language with

- for each $i \in I$, n many predicate symbols, $P_1^i, \ldots, P_n^i$, each represents a cell in the n-partition of Γ_i
- for $1 \leq j \leq m$, introduce constant symbols, $c_1^j, \ldots, c_{k_j}^j$, each name the corresponding element of $\mathcal{A}_j$

For each $i \in I$, let

$$\Theta_i = \forall x(\bigvee_{1 \leq h \leq n} P_h^i x)$$

and

$$\Omega_i = \Big\{\Big(\bigwedge_{1 \leq h \leq n} (\bigvee_{1 \leq j \leq k_1} \neg P_h^i c_j^1)\Big) \wedge \ldots \wedge \Big(\bigwedge_{1 \leq h \leq n} (\bigvee_{1 \leq j \leq k_m} \neg P_h^i c_j^m)\Big) \wedge \Theta_i\Big\}$$

Let

$$\Omega = \bigcup_{i \in I} \Omega_i$$

We obtain our first order theory T by adding all elements of Ω as proper axioms to first order logic. But our assumption is that Σ is not an n-trace over Γ_i, for each $i \in I$, so every finite subset of Ω has a model. By first order compactness, T has a model and thus Ω has a model. Hence there must be an n-partition of Γ such that no cell of the partition contains any element of Σ, i. e. Σ is not an n-trace over Γ. ∎

Theorem 6.3.1

Trace Compactness: let Γ be a non-empty set and Σ be a collection of finite subsets of Γ. Then Σ is an n-trace over Γ iff there is a $\Gamma_0 \subseteq_{fin} \Gamma$ and a $\Sigma_0 \subseteq_{fin} \Sigma$ such that Σ_0 is an n-trace over Γ_0.

Proof:

($\Leftarrow$): Assume that Σ_0 is an n-trace over Γ_0, where $\Gamma_0 \subseteq_{fin} \Gamma$ and $\Sigma_0 \subseteq_{fin} \Sigma$. To show that Σ is an n-trace over Γ, it suffices to show that Σ_0 is an n-trace over Γ. Clearly each n-partition of Γ must also partition Γ_0 into n or fewer cells. If an n-partition of Γ partitions Γ_0 into n cells, then by our initial hypothesis, some element of Σ_0 is a subset of some cell in the partition. And if an n-partition of Γ partitions Γ_0 into m cells, where $m < n$, then by lemma (6.3.1), Σ_0 is an m-trace over Γ_0, so some element of Σ_0 must also be a subset of some cell of the partition. So either way, some element of Σ_0 is a subset of some cell in each n-partition of Γ. Hence Σ_0 is an n-trace over Γ.

($\Rightarrow$): Assume that Σ is an n-trace over Γ. By lemma (6.3.2), $\exists \Sigma_0 \subseteq_{fin} \Gamma$ such that Σ_0 is an n-trace over Γ. By lemma (6.3.3), $\exists \Gamma_0 \subseteq_{fin} \Gamma$ such that Σ_0 is an n-trace over Γ_0. ∎

Having now established the compactness of traces, we are now in a position to prove the compactness of coherence level of sets.

Theorem 6.3.2
Level Compactness: For $\Gamma \subseteq \Phi$, if $\ell(\Gamma) > n$, then there is a finite subset Γ' of Γ, such that $\ell(\Gamma') > n$.

Proof:
Let $\Gamma \subseteq \Phi$ such that $\ell(\Gamma) > n$. Let

$$\Sigma = \{\mathcal{A} : \mathcal{A} \subseteq_{fin} \Gamma \text{ and } \mathcal{A} \vdash_{\wedge} \bot\}$$

By our assumption that $\ell(\Gamma) > n$, Σ is an n-trace over Γ. By theorem (6.3.1), $\exists \Gamma_0 \subseteq_{fin} \Gamma$, $\exists \Sigma_0 \subseteq_{fin} \Sigma$ such that Σ_0 is an n-trace over Γ_0. But then every n-partition of Γ_0 must contain a cell which includes some element of Σ_0. But $\Sigma_0 \subseteq \Sigma$. Hence every n-partition of Γ_0 contains an inconsistent cell, i.e. $\ell(\Gamma_0) > n$ as required. ∎

Theorem 6.3.3
Compactness of $\vdash_n$: For $\Gamma \subseteq \Phi$, $A \in \Phi$, if $\Gamma \vdash_n A$, then $\exists \Gamma_0 \subseteq_{fin} \Gamma$ such that $\Gamma_0 \vdash_n A$.

Proof:
Let $\Gamma \subseteq \Phi$, $A \in \Phi$, such that $\Gamma \vdash_n A$. Then by proposition (6.3.1), $\ell([\wedge\neg A]\Gamma) > n$. By theorem (6.3.2), there is $\Gamma_0 \subseteq_{fin} \Gamma$ such that $\ell([\wedge\neg A]\Gamma_0) > n$. Hence by proposition (6.3.1) again, $\Gamma_0 \vdash_n A$. ∎

6.4 Completeness of n-Forcing

In [11], the completeness of K_n is proven in two stages: the first stage is to show the compactness of $\vdash_n$ using compactness of hypergraph colouring, the second stage is to show the completeness of $\vdash_n$ by presenting an axiomatisation of $\vdash_n$ in term of Gentzen-style sequents. In this section we will present the Apostoli-Brown axiomatisation of $\vdash_n$.

Axiom Schema:

$$\frac{}{A \Vdash_n A}\,[\mathrm{Ref}]$$

Rules Schemata:

$$\frac{\Gamma \Vdash_n A}{\Gamma, \Delta \Vdash_n A}\,[\mathrm{Mon}] \qquad \frac{\Gamma \Vdash_n A_1 \cdots \Gamma \Vdash_n A_{n+1}}{\Gamma \Vdash_n \bigvee_{1 \leq i < j \leq n+1} (A_i \wedge A_j)}\,[\mathrm{RK}_n]$$

$$\frac{\Gamma \Vdash_n A \qquad \Gamma, A \Vdash_n B}{\Gamma \Vdash_n B}\,[\mathrm{N\ Cut}] \qquad \frac{\Gamma \Vdash_n A \qquad \vdash_\wedge B}{\Gamma \Vdash_n B}\,[\wedge\ \mathrm{Cut}]$$

Given the usual notions of *derivation of a sequent* and *derivable sequent*, we obtain the following result:

Theorem 6.4.1
Completeness for n-forcing (Apostoli and Brown): Let $\Gamma \subseteq \Phi$, $A \in \Phi$. Then

$$\Gamma \vdash_n A \Longleftrightarrow \Gamma \Vdash_n A$$

Proof:
see proof of theorem (5.2), p.839, in [11] ∎

6.5 Completeness of K_n^m

We are now in a position to show that just as the schema $[K_n]$ yields closure under the rule schema $[RT_n]$, the schema $[K_n^m]$ yields closure under the rule schema

$$\frac{\Gamma \vdash_n A}{C_B[\Gamma] \vdash_\wedge C_B[A]}\,[\mathrm{RT}_n^m]$$

where $C_B[\Gamma] = \{C_B[D] : D \in \Gamma\}$ for for some $C = \Box(C_1, \ldots, C_{m-1}, B)$.

Theorem 6.5.1
Let $\Gamma \subseteq \Phi$, $A \in \Phi$ *and* $C = \Box(C_1, \ldots, C_{m-1}, B)$*, then*

$$\Gamma \Vdash_n A \implies C_B[\Gamma] \vdash_\wedge C_B[A]$$

Proof:
The proof is by induction on the complexity of the derivation of $\Gamma \Vdash_n A$.
Basis: $\Gamma \Vdash_n A$ is an axiom. Then $\Gamma = \{A\}$. But $C_B[A] \vdash_\wedge C_B[A]$. So $C_B[\Gamma] \vdash_\wedge C_B[A]$.

Induction Step: Assume that $\Gamma \Vdash_n A$ is obtained as an endsequent by application of one of the rule schemata, and that the theorem holds with respect to all proper subderivations therein. There are four cases:

(1) [Mon]: the result is obtained by the induction hypothesis and the monotonicity of $\vdash_\wedge$.

(2) [$\wedge$ Cut]: use the induction hypothesis, PL and [RM_n^m].

(3) [N Cut]: use the induction hypothesis and PL.

(4) [RK_n^m]: Then

$$A = \bigvee_{1 \leq i < j \leq n+1} (A_i \wedge A_j)$$

such that $\Gamma \Vdash_n A_k$, for each $k \leq n+1$. By the induction hypothesis, $C_B[\Gamma] \vdash_\wedge C_B[A_k]$, for each $k \leq n+1$. So

$$C_B[\Gamma] \vdash_\wedge C_B[A_1] \wedge \ldots \wedge C_B[A_{n+1}]$$

But

$$C_B[A_1] \wedge \ldots \wedge C_B[A_{n+1}] \to C_B[\bigvee_{1 \leq i < j \leq n+1} (A_i \wedge A_j)]$$

is an instance of [K_n^m], hence by PL,

$$C_B[\Gamma] \vdash_\wedge C_B[\bigvee_{1 \leq i < j \leq n+1} (A_i \wedge A_j)]$$

follows immediately. This completes the induction. ∎

We now give the canonical model construction for K_n^m, and show that every non-theorem of K^m_n is falsified in the corresponding canonical model.

Definition 6.5.1

The Canonical Model of K^m_n, $\mathfrak{M}^\Lambda = \langle \mathcal{U}^\Lambda, \mathcal{R}^\Lambda, \mathcal{V}^\Lambda \rangle$ *where*

- $\mathcal{U}^\Lambda = \{x : x$ *is a maximal* K^m_n *consistent set*$\}$
- $\mathcal{R}^\Lambda \subseteq \wp(\mathcal{U}^\Lambda)^{m-1} \times \mathcal{U}^{\Lambda^{n+1}}$ *is defined by*
  $\forall x y_1, \ldots, y_n \in \mathcal{U}, \forall \mathcal{A}_1, \ldots, \mathcal{A}_{m-1} \subseteq \mathcal{U}, \forall A_1, \ldots, A_{m-1}, B \in \Phi,$

  $\mathcal{R}^\Lambda \mathcal{A}_1, \ldots, \mathcal{A}_{m-1}, x y_1, \ldots, y_n \Longleftrightarrow$
  $\forall i_{(1 \leq i \leq m-1)} \mathcal{A}_i = |A_i|^\Lambda$ *and* $\Box(A_1, \ldots, A_{m-1}, B) \in x \Longrightarrow \exists j_{(1 \leq j \leq n)} : B \in y_j$

  where $C \in \Phi$, $|C|^\Lambda = \{x \in \mathcal{U}^\Lambda : C \in x\}$ *(note that* $|C|^\Lambda = |D|^\Lambda$ *iff* $\vdash_\Lambda C \leftrightarrow D$*).*
- $\mathcal{V}^\Lambda : At \longrightarrow \wp(\mathcal{U}^\Lambda)$ *is defined by*

$$\forall x \in \mathcal{U}^\Lambda,\ \forall p_i \in At,\ x \in \mathcal{V}^\Lambda(p_i) \Longleftrightarrow p_i \in x$$

Theorem 6.5.2

The Fundamental Theorem for K^m_n *modal logics:* $\forall x \in \mathcal{U}^\Lambda$, $\forall A \in \Phi$,

$$\models^{\mathfrak{M}}_x A \Longleftrightarrow A \in x$$

Proof:

The proof is by induction on the complexity of A. The basis is given by the definition of $\mathcal{V}^\Lambda$. For the induction step, we assume the induction hypothesis that the theorem holds with respect to all sub-formulae of A and show that the theorem holds for A. The cases for the truth functional connectives are trivial, we'll consider the case where $A = \Box(C_1, \ldots, C_{m-1}, B)$. We need to show that

$$\models^{\mathfrak{M}}_x \Box(C_1, \ldots, C_{m-1}, B) \Longleftrightarrow \Box(C_1, \ldots, C_{m-1}, B) \in x$$

($\Leftarrow$): follows immediately from the definition of $\mathcal{R}^\Lambda$ and the induction hypothesis that $\|C_i\|^{\mathfrak{M}^\Lambda} = |C_i|^\Lambda$ for $i \leq m-1$.

($\Rightarrow$): Assume that $\Box(C_1, \ldots, C_{m-1}, B) \notin x$. By the induction hypothesis, it suffices to show that $\exists y_1, \ldots, y_n \in \mathcal{U}^\Lambda$:

$$\mathcal{R}^\Lambda |C_1|^\Lambda, \ldots, |C_{m-1}|^\Lambda, x y_1, \ldots, y_n \text{ and } \forall j_{(1 \leq j \leq n)} \neg B \in y_j$$

We will construct such $y_1, \ldots, y_n \in \mathcal{U}^\Lambda$. Let

$$\Box(x) = \{D : \Box(C_1, \ldots, C_{m-1}, D) \in x\}$$

Claim: $\ell(\Lambda \neg B[\Box(x)]) \leq n$

Proof of Claim: Towards a contradiction we assume that $\ell(\Lambda \neg B[\Box(x)]) > n$. Then by proposition (6.3.1), $\Box(x) \vdash_n B$. By the rule $[RT_n^m]$, $x \vdash_\Lambda \Box(C_1, \ldots, C_{m-1}, B)$. By the deductive closure of x, $\Box(C_1, \ldots, C_{m-1}, B) \in x$ which contradicts our initial assumption. This completes the proof of our claim.

By our claim, $\exists \pi \in \prod_n(\Box(x))$ such that $\pi = \{\mathcal{C}_1, \ldots, \mathcal{C}_n\}$ and for $i \leq n$, $\mathcal{C}_i \cup \{\neg B\}$ is Λ-consistent. By Lindenbaum's lemma, for each $i \leq n$, $\mathcal{C}_i \cup \{\neg B\}$ can be extended to $y_i \in \mathcal{U}^\Lambda$. It remains to be proven that

$$\mathcal{R}^\Lambda |C_1|^\Lambda, \ldots, |C_{m-1}|^\Lambda, x y_1, \ldots, y_n$$

Let $C_1, \ldots, C_{m-1}, D \in \Phi$ be such that $|C_j|^\Lambda = |A_j|^\Lambda$ for every $j \leq m-1$ and $\Box(C_1, \ldots, C_{m-1}, D) \in x$. Then clearly, for each $j \leq m-1$, $\vdash_\Lambda C_j \leftrightarrow A_j$ and $\vdash_\Lambda D \rightarrow D$. Hence by $[RM_n^m]$,

$$\vdash_\Lambda \Box(C_1, \ldots, C_{m-1}, D) \rightarrow \Box(A_1, \ldots, A_{m-1}, D)$$

But $\Box(C_1, \ldots, C_{m-1}, D) \in x$, so $\Box(A_1, \ldots, A_{m-1}, D) \in x$. Hence, $D \in \Box(x)$. But by our construction $\Box(x) \subseteq \bigcup_{1 \leq i \leq n} y_i$. Hence, $D \in \bigcup_{1 \leq i \leq n} y_i$ as required. This completes the inductive proof of the fundamental theorem. ∎

6.6 Further Work

It is clear that for each fixed m, we have a descending sequence of K_n^m logics ordered by inclusion:

$$K_1^m \supseteq K_2^m \supseteq K_3^m \supseteq \ldots$$

In [100], Jennings and Schotch show that the limit of the descending sequence of K_n logics is finitely axiomatisable by PL, [N] and [RM] alone. To show that this is indeed the case, Jennings and Schotch show that

1. N is determined by a class, $\mathfrak{C}$, of locale frames.

2. Any formula that fails in $\mathfrak{C}$ also fails in the class of relational frames of rank $n+1$ where $n+1$ is the arity of the relation.

Since K_n is determined by the class of relational frames of rank $n+1$, it follows immediately from 1. and 2. that every non-theorem of N is a non-theorem of K_n and hence $\bigcap_{n<\omega} K_n = N$.

Question 6.6.1
For each fixed m*, does* $\bigcap_{n<\omega} K_n^m = N^m$*, where* N^m *is axiomatised by* PL $[N^m]$ *and* $[RM_n^m]$*?*

Chapter 7

Hypergraph Satisfiability

7.1 Introduction

In [112], Kolany introduces the notion of weak satisfiability on hypergraphs, a generalisation of Cowen's notion of strong satisfiability on hypergraphs [51], and shows that the compactness property of weak satisfiability on hypergraphs is, in ZF set theory, equivalent to BPI, i.e. to the statement that every Boolean Algebra contains an ultrafilter [50]. Kolany's notion of weak satisfiability provides a graph theoretic representation of a wide range of combinatorial problems, including the satisfiability of propositional formulae. In this chapter, we'll generalise Kolany's idea and introduce the notion of n-satisfiability on hypergraphs. The motivation for such a generalisation originates in the works [101; 164; 167]. In their study of inconsistency-tolerant logic, Jennings and Schotch observe that a set of unsatisfiable formulae may be partitioned into subsets which are satisfiable individually. Consider for instance,

$$\Sigma = \{p \wedge q,\ \neg p \wedge \neg q,\ p \wedge \neg q,\ r\}$$

Although Σ is unsatisfiable, it can be partitioned into 3 satisfiable subsets but every partition of Σ into 2 subsets contains at least one unsatisfiable subset. Thus we may think of Σ as a 3-satisfiable set but not a 2-satisfiable set. More precisely, for a set of formulae, Σ, we say that it is n-satisfiable iff there is a partition of Σ into n or fewer satisfiable subsets. For a set of satisfiable formulae, we may conveniently treat it as a 1-satisfiable set. The *incoherence level* of a set Σ, $\ell(\Sigma)$, is then defined as the least n such that Σ is n-satisfiable, if there is no such n then $\ell(\Sigma) = \infty$; if $\Sigma = \emptyset$ or is satisfiable, we let $\ell(\Sigma) = 1$. The observations of Jennings and Schotch form the basis for a species of paraconsistent logics in which the classical rule of *ex contradictione quodlibet* is not derivable; i.e. $A, \neg A \not\vdash B$. Thus the Jennings-Schotch's logics provide one particular strategy for reasoning with inconsistent information. More interestingly however, the Jennings-Schotch's logics have startling connections with modal logics as well as the infamous four colour theorem, i.e. every planar graph is 4-colourable. We cannot give a full summary of their works here. The reader is advised to consult [9; 11; 44] for details.

Our aim here is to develop hypergraph theoretic counterparts to the notion of n-satisfiability and the related notion of incoherence level of a set. Our motivation is partly to continue the theoretical exploration of the connection between paraconsistent reasoning and hypergraph theory. But we are also interested in developing a general framework for visualising logical problems that involve reasoning with inconsistent information. In this chapter, we'll develop a general notion of n-satisfiability on hypergraphs which subsumes Kolany's notion of weak satisfiability. We'll also show that the compactness statement for n-satisfiability on hypergraphs is equivalent to BPI in ZF set theory. We give a syntactic characterisation of n-satisfiability on hypergraphs in terms of a resolution style proof procedure. A general notion of consequence relation based on hypergraphs will also be introduced. We'll conclude with a discussion of a conjecture of Cowen relating BPI and complexity theory.

First we recall Kolany's definitions.

Definition 7.1.1
A hypergraph H is a pair (V, E) where V is a non-empty set of vertices or literals, and its finite subsets are called clauses; E is a collection of non-empty subsets of V. The elements of E are called edges. A hypergraph is compact if all edges are finite. A hypergraph is a graph if $\forall e \in E,\ |e| = 2$.

Let $n \geq 2$. A (vertex) n-colouring of $H = (V, E)$ is a function $c : V \longrightarrow \{1, \ldots, n\}$ such that all edges are non-monochromatic under c. We say that H is n-colourable iff there is an n-colouring of H. The chromatic number, $\chi(H)$, of H is the least n such that H is n-colourable.

We note that if H contains a singleton edge, then H is not n-colourable. We are concerned only with compact hypergraphs without singleton edges here. Kolany's notion of weak satisfiability is defined as follows:

Definition 7.1.2
Let H be an arbitrary but fixed hypergraph. Let Γ be a set of literals on H and Σ be a family of clauses on H. Then Γ weakly satisfies Σ (on H) iff:

1. $\forall e \in E$, $e \not\subset \Gamma$ *(Γ is consistent)*

2. $\forall \sigma \in \Sigma, \Gamma \cap \sigma \neq \emptyset$ *(Γ pierces Σ)*

Σ is weakly satisfiable on H iff some $\Gamma \subseteq V$ weakly satisfies Σ.

It is very easy to visualize Kolany's notion of weak satisfiability on hypergraphs. In fact the problem of determining whether a family of clauses is weakly satisfiable on a given hypergraph is structurally similar to the travelling salesman problem (see [76]). We forgo formal definitions here and adopt a more intuitive presentation using descriptions such as points, regions and tours. In figure (7.1), vertices are represented by black dots and edges are represented by blue lines or blue regions. A family of clauses, Σ, is represented by red lines or red regions. To determine whether Σ is weakly satisfiable is simply a matter of finding a *tour* that

1. passes each red region at least once and

2. avoids passing through every point in any given blue region.

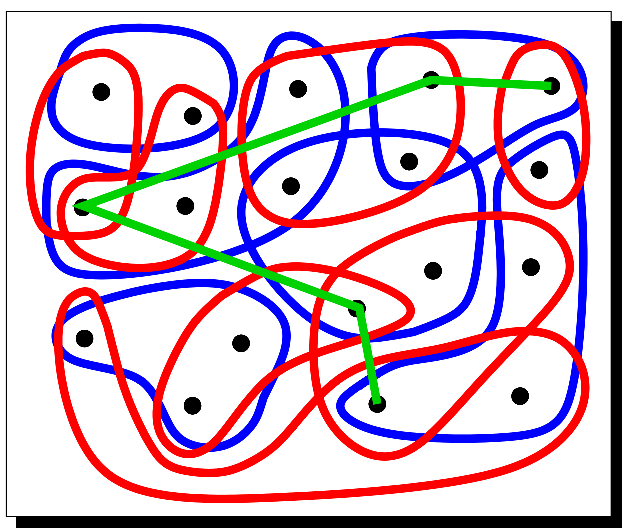

Figure 7.1: n-satisfiability on hypergraphs.

The green line represents a tour that weakly satisfies Σ. In [112], Kolany proved the following:

Theorem 7.1.1
Kolany [112]
(1) Let H be a compact hypergraph and Σ a set of clauses on H. Then Σ is weakly satisfiable iff every finite subset of Σ is weakly satisfiable.

(2) The compactness theorem (in (1)) for weak satisfiability on hypergraphs is equivalent to BPI *(in* ZF *set theory).*

7.2 n-satisfiability on Hypergraphs

To represent the Jennings-Schotch notion of n-satisfiability via hypergraphs, we need to first extend Kolany's definition of weak satisfiability. In particular a propositional formula can be represented by a finite set of clauses, i.e. a finite set of finite sets of vertices, and a set of propositional formulae can be represented by a collection of finite sets of clauses. We call a finite set of clauses a *formula* of a hypergraph. Formulae will be denoted by A, B, C, ...etc. We now introduce a generalised version of Kolany's notion of weak satisfiability.

Definition 7.2.1
Let Σ be a set of formulae of a hypergraph $H = (V, E)$ and $\Gamma_1, \dots, \Gamma_n \subseteq V$. Then $\Gamma_1, \dots, \Gamma_n$ n-satisfy Σ iff

1. $\forall e \in E, \forall i \leq n, e \not\subset \Gamma_i$ *(Γ_i is consistent)*

2. $\forall A \in \Sigma, \exists i \leq n$: $\forall \sigma \in A, \Gamma_i \cap \sigma \neq \emptyset$ *(each $A \in \Sigma$ is weakly satisfied or covered by some Γ_i)*

We say that Σ is n-satisfiable on H iff some $\Gamma_1, \dots, \Gamma_n$ n-satisfy Σ on H.

Remark 7.2.1 We note that 1-satisfiability of a set of formulae Σ on H is equivalent to the weak satisfiability of the set of clauses $\bigcup \Sigma$. Conversely, if Σ is a collection of clauses, then the weak satisfiability of Σ on H is equivalent to the 1-satisfiability of $\{\Sigma' : \Sigma' \subseteq_{fin} \Sigma\}$.

Notice that in our definition we do not require that $\Gamma_1, \dots, \Gamma_n$ be distinct, so our definition says that a set of formulae is n-satisfiable on H iff they are covered by n or fewer consistent sets of vertices. To illustrate, consider the following example:

Example 7.2.1

$$V = \{u, v, w, x, y, z\} \qquad E = \Big\{\{u, v\},\ \{w, x\},\ \{y, z\}\Big\}$$

$$\Sigma = \Big\{\big\{\{u, w\},\ \{z\}\big\},\ \big\{\{v\},\ \{x\},\ \{w, y\}\big\},\ \big\{\{u, w\},\ \{y\}\big\}\Big\}$$

It is straightforward to verify that Σ is neither 1-satisfiable nor 2-satisfiable; but it is 3-satisfiable, e.g. $\{u, z\}$, $\{v, x, y\}$ and $\{u, y\}$ 3-satisfy Σ. Alternatively, we may view our example as a propositional graph with $u = p$, $v = \neg p$, $w = q$, $x = \neg q$, $y = r$ and $z = \neg r$, in which case rewriting every member of Σ in conjunctive normal form, we have:

$$\Sigma' = \big\{(p \vee q) \wedge \neg r,\ \neg p \wedge \neg q \wedge (q \vee r),\ (p \vee q) \wedge r\big\}$$

Clearly, $\ell(\Sigma') = 3$. In terms of visual representation, we can treat the n-satisfiability problem on hypergraphs as a *multi-dimensional* version of the weak satisfiability problem on hypergraphs. The dimension is given by $|\Sigma|$ where each $A \in \Sigma$ is a distinct dimension *above* the hypergraph H. To determine whether Σ is n-satisfiable on H is then to find n or fewer tours such that

1. every region from a given dimension is visited by one of the tours, and
2. no one tour passes all vertices of any given edge.

For a set of formulae, Σ, on a hypergraph H, we can define a function analogous to the incoherence level of Jennings and Schotch.

Definition 7.2.2
Let $\infty \notin \mathbb{N}$. Then relative to a hypergraph H, the λ_H-level of Σ is defined as follows:

$$\lambda_H(\Sigma) = \begin{cases} \text{the least } n \text{ such that} & \text{if } n \in \mathbb{N} \text{ exists} \\ \Sigma \text{ is } n\text{-satisfiable} & \\ & \\ \infty & \text{otherwise} \end{cases}$$

In effect, $\lambda_H(\Sigma)$ computes the least number of tours required to achieve (1) and (2) above for Σ. Moreover, our notion of n-satisfiability on hypergraphs is also closely related to the ordinary notion of n-colourability of hypergraphs. The λ_H value of a set Σ is in fact equal to the proper chromatic number of an appropriate hypergraph.

Theorem 7.2.1
Let Σ be a set of formulae of a hypergraph H. Let $H_\Sigma = (V_\Sigma, E_\Sigma)$ be the hypergraph with $V_\Sigma = \Sigma$ and

$$E_\Sigma = \big\{\Sigma' \subseteq_{fin} \Sigma :\ \Sigma' \neq \emptyset \ and\ \lambda_H(\Sigma') > 1\big\}$$

Then $\lambda_H(\Sigma) = \chi(H_\Sigma)$.

Proof:

Let H, Σ, and H_Σ be defined as above. Let $\chi(H_\Sigma) = n$. We show that $\lambda_H(\Sigma) \leq n$. Let c be a proper n-colouring of H_Σ. For each $i \leq n$, let $\sigma_i = \{A \in \Sigma : \ c(A) = i\}$. By remark (7.2.1) and the compactness of weak satisfiability (see [112] p.396), it is straightforward to verify that σ_i is 1-satisfiable on H. Let Γ_i 1-satisfy σ_i on H. To show $\lambda_H(\Sigma) \leq n$, it suffices to show that $\Gamma_1, \ldots, \Gamma_n$ n-satisfy Σ, i.e. each $A \in \Sigma$ is covered by some Γ_i:

$$\begin{aligned} A \in \Sigma &\implies A \in \sigma_i \text{ for some } i \leq n \\ &\implies \Gamma_i \text{ 1-satisfies } \sigma_i \\ &\implies \Gamma_i \text{ 1-satisfies } A \\ &\implies \forall \alpha \in A,\ \Gamma_i \cap \alpha \neq \emptyset \end{aligned}$$

Towards a contradiction, let $\lambda_H(\Sigma) = k$ where $k < n$. Let $\Gamma_1, \ldots, \Gamma_k$ k-satisfy Σ. For each $i \leq k$, let $\gamma_i = \Gamma_i \cap \bigcup\bigcup \Sigma$. Then $\gamma_1, \ldots, \gamma_k$ also k-satisfy Σ. To show that there is a proper k-colouring on H_Σ, we first define the following sequence of sets:

$$\begin{aligned} \Sigma_1 &= \{A \in \Sigma : \ \forall \alpha \in A,\ \alpha \cap \gamma_1 \neq \emptyset\} \\ \Sigma_2 &= \{A \in \Sigma \setminus \Sigma_1 : \ \forall \alpha \in A,\ \alpha \cap \gamma_2 \neq \emptyset\} \\ &\ \vdots \\ \Sigma_k &= \{A \in \Sigma \setminus \Sigma_{k-1} : \ \forall \alpha \in A,\ \alpha \cap \gamma_k \neq \emptyset\} \end{aligned}$$

Clearly, $\Sigma_i \cap \Sigma_j = \emptyset$, for each $i \neq j \leq k$. Moreover, Σ_i is 1-satisfied by γ_i and thus no edge of H_Σ is a subset of Σ_i. Now define $c : \Sigma \longrightarrow \{1, \ldots,\ k\}$ such that $c(A) = i$ iff $A \in \Sigma_i$. Then c is a proper k-colouring of H_Σ. This contradicts our assumption that $\chi(H_\Sigma) = n$. Hence $\lambda_H(\Sigma) = \chi(H_\Sigma)$ as required. ∎

It is interesting to note that proper n-colouring problems can also be restated as n-satisfiability problems. Let $H = (V, E)$ be a hypergraph such that $V = \{x,\ y,\ z,\ \ldots\}$, then the n-satisfiability of $\{\{\{x\}\},\ \{\{y\}\},\ \{\{z\}\},\ \ldots\}$ on H is equivalent to the proper n-colourability of H. We can in fact establish the compactness of n-satisfiability directly from the compactness of proper n-colourability of hypergraphs.

Theorem 7.2.2

A compact hypergraph H is properly n-colourable iff every finite sub-hypergraph of H is properly n-colourable.

Proof:
Our proof of trace compactness from the previous chapter suffices (theorem (6.3.1)). We note however that the Axiom of Choice is not required in the proof. ∎

Theorem 7.2.3
Let H be a compact hypergraph and Σ be a set of formulae of H. Then Σ is n-satisfiable iff every finite subset of Σ is n-satisfiable.

Proof:
One direction is trivial. For the other direction, we assume that Σ is not n-satisfiable on H. By theorem (7.2.1), $\chi(H_\Sigma) > n$ and thus H_Σ is not properly n-colourable. By theorem (7.2.2), some finite subhypergraph $H'_\Sigma = (V_{H'_\Sigma}, E_{H'_\Sigma})$ of H_Σ is not properly n-colourable. Clearly, $V_{H'_\Sigma} \subseteq_{fin} \Sigma$ and $V_{H'_\Sigma}$ is not n-satisfiable on H. ∎

Corollary 7.2.1
Let H be a compact hypergraph and Σ be a set of formulae of H. Then $\lambda_H(\Sigma) \leq n$ iff for every finite subset Σ' of Σ, $\lambda_H(\Sigma') \leq n$.

Theorem 7.2.4
The compactness theorem for n-satisfiability on hypergraphs is equivalent to BPI in ZF set theory.

Proof:
Since BPI is equivalent to the compactness theorem for first order logic in ZF (see [20] p.104), it follows immediately from the proof of theorem (6.3.1) that BPI implies the compactness of n-satisfiability on hypergraphs. For the converse, Kolany showed that the compactness of weak satisfiability is equivalent to BPI. But the compactness of weak satisfiability is just a special case of the compactness of n-satisfiability with $n = 1$ (see remark (7.2.1)). Hence the desired result follows immediately. ∎

In [112], Kolany gives a partial list of problems in different branches of mathematics that are equivalent to the weak satisfiability of some family of clauses on an appropriate hypergraph. Many compactness statements in different areas of mathematics can in fact be viewed as instances of compactness of weak satisfiability on hypergraphs. They include:

satisfiability of propositional formulae

Axiom of Choice for finite sets

let Δ be a collection of pairwise disjoint non-empty finite sets; then f is a choice function on Δ if $f : \Delta \longrightarrow \bigcup \Delta$ and for all $\delta \in \Delta$, $f(\delta) \in \delta$.

R-consistent choices (equivalent to BPI in ZF)

let Δ be a collection of pairwise disjoint non-empty finite sets and let $\mathcal{R}$ be a symmetric binary relation on $\bigcup \Delta$; then Π is an $\mathcal{R}$-consistent choice if for all $\alpha, \beta \in \Pi$, $\alpha \mathcal{R} \beta$.

graph n-colourability

polynomial equation system solvability over a finite field

(equivalent to BPI in ZF)

systems of distinct representatives

let Δ be a collection of non-empty finite sets, then an injective function f is a System of Distinct Representatives for Δ if $f : \Delta \longrightarrow \bigcup \Delta$ and for all $\delta \in \Delta$, $f(\delta) \in \delta$.

Remark (7.2.1) and theorem (7.2.4) make it clear that any such problem and its corresponding compactness statement can be restated as an n-satisfiability problem and a corresponding compactness statement on appropriate hypergraphs.

7.3 Resolution and n-satisfiability

The notion of n-satisfiability is essentially a semantic notion. We now wish to consider a purely syntactic characterisation of the same notion. The characterization given here is analogous to the resolution rule that J. A. Robinson developed for automated theorem proving.

Definition 7.3.1

Let $\alpha_1, \ldots, \alpha_m$ be a set of clauses on an arbitrary but fixed hypergraph H and let $e = \{a_1, \ldots, a_m\}$ be an edge of H. Then the set

$$\alpha = \bigcup_{i=1}^{m} (\alpha_i \setminus \{a_i\})$$

is said to result by resolution on e if $a_i \in \alpha_i$ for $1 \leq i \leq m$. In which case, we write $\alpha_1, \ldots, \alpha_m \vdash_e \alpha$. If A is a set of clauses on H, the closure of A under the resolution rule on edges of H is denoted by $[A]_H$.

Consider again example (7.2.1). It is easy to show that $\emptyset \in [\bigcup \Sigma]_H$; we display two resolution proofs:

$$\{u,w\},\{v\} \vdash_{\{u,v\}} \{w\} \qquad \{u,w\},\{x\} \vdash_{\{w,x\}} \{u\}$$
$$\{w\},\{x\} \vdash_{\{w,x\}} \emptyset \qquad \{u\},\{v\} \vdash_{\{u,v\}} \emptyset$$

In terms of weak satisfiability on hypergraphs, the resolution rule is a sound and complete rule.

Theorem 7.3.1
(Kolany) Let H *be a compact hypergraph and* A *be a finite collection of clauses. Then* $\emptyset \in [A]_H$ *iff* A *is not weakly satisfiable on* H.

For a set of formulae Σ on H, we call a function $f : \Sigma \longrightarrow n$ a Σ-n-colouring. Where $i \leq n$, we take $f^{-1}[i] = \{A \in \Sigma : \ f(A) = i\}$. We can now characterise n-satisfiability in terms of resolution and Σ-n-colouring.

Theorem 7.3.2
Let Σ *be a finite set of formulae on a compact hypergraph* H. *Then* Σ *is not* n*-satisfiable iff for every* Σ*-*n*-colouring* f, $\exists i \leq n : \ \emptyset \in [\cup(f^{-1}[i])]_H$.

Proof:
$(\Rightarrow)$ Let f be a Σ-n-colouring such that $\forall i \leq n$, $\emptyset \notin [\cup(f^{-1}[i])]_H$. Then by theorem (7.3.1), for each $i \leq n$, $\cup(f^{-1}[i])$ is weakly satisfiable on H. Let Γ_i weakly satisfy $\cup(f^{-1}[i])$ on H for each $i \leq n$. Then $\Gamma_1, \ldots, \Gamma_n$ n-satisfy Σ on H.

$(\Leftarrow)$: Assume that for any Σ-n-colouring f, $\exists i \leq n : \ \emptyset \in [\cup(f^{-1}[i])]_H$. Then by theorem (7.3.1), $\exists i \leq n : \ \cup(f^{-1}[i])$ is not weakly satisfiable on H. Towards a contradiction, let $\Gamma_1, \ldots, \Gamma_n$ n-satisfy Σ on H. Define the Σ-n-colouring g, such that

$$\begin{aligned}
g^{-1}[1] &= \{A \in \Sigma : \ \forall \alpha \in A, \ \alpha \cap \Gamma_1 \neq \emptyset\} \\
g^{-1}[2] &= \{A \in \Sigma \setminus g^{-1}[1] : \ \forall \alpha \in A, \ \alpha \cap \Gamma_2 \neq \emptyset\} \\
&\vdots \\
g^{-1}[n] &= \{A \in \Sigma \setminus g^{-1}[n-1] : \ \forall \alpha \in A, \ \alpha \cap \Gamma_n \neq \emptyset\}
\end{aligned}$$

Since g is a Σ-n-colouring on Σ, $\exists i \leq n : \ \cup(g^{-1}[i])$ is not weakly satisfiable on H. This implies that Γ_i does not weakly satisfy $\cup(g^{-1}[i])$ and thus $\Gamma_1, \ldots, \Gamma_n$ fail to n-satisfy Σ.

∎

Combining theorem (7.2.3) and theorem (7.3.2), we can immediately derive the following theorem:

Theorem 7.3.3

For a set of formulae Σ *on a compact hypergraph* H*,* Σ *is* n*-satisfiable iff there exists a* Σ*-*n*-colouring* f *such that* $\forall i \leq n,\ \emptyset \notin [\cup(f^{-1}[i])]_H$.

In terms of complexity, the decision version of our problem is **NP**-complete.

Theorem 7.3.4

The decision problem for determining whether a finite set of formulae Σ *is* n*-satisfiable on a compact hypergraph* H *(*H*-*n*-**SAT**) is **NP**-complete.*

Proof:

It is easy to see that H-n-**SAT** is at least **NP**-hard since an instance of our problem is just **SAT** which is **NP**-complete. It is also clear that H-n-**SAT**$\in$**NP**, since a nondeterministic turing machine can guess a sequence $(\Gamma_1, \ldots, \Gamma_n)$, each of which is a subset of Σ, and verify in polynomial time whether $(\Gamma_1, \ldots, \Gamma_n)$ n-satisfy Σ. ∎

7.4 n-Consequence Relations

We have seen that the notion of weak satisfiability can be generalised nicely to the notion of n-satisfiability. In this section, we'll develop two notions of consequence relations based on n-satisfiability on hypergraphs. Again, these are natural generalisations of Kolany's notion of consequence operations based on weak satisfiability (see [113]).

In the subsequent exposition, we let $H = (V, E)$ be a fixed compact hypergraph without singleton edges. The set of all formulae on H will be denoted by Φ. We'll use the usual Σ, A and Σ, Δ to denote $\Sigma \cup \{A\}$ and $\Sigma \cup \Delta$ respectively.

Definition 7.4.1

$\Gamma_1, \ldots, \Gamma_n \subseteq V$ *is an* n*-model of* Σ *on* H *iff* $\Gamma_1, \ldots, \Gamma_n$ n*-satisfy* Σ *on* H*. If in addition,* $\Gamma_1, \ldots, \Gamma_n$ n*-cover* E*, i.e.* $\forall e \in E, \exists i \leq n$*:* $\Gamma_i \cap e \neq \emptyset$*, then* $\Gamma_1, \ldots, \Gamma_n$ *is an* n^+*-model of* Σ *on* H*.* $\models_n$ *and* $\models_n^+$ *are defined as follows:*

- $\Sigma \models_n \{A\}$ *iff every* n*-model of* Σ *is an* n*-model of* $\{A\}$ *(on* H*). The set of* n*-consequences of* Σ*,* $\mathbf{C}_n(\Sigma) = \{A : \Sigma \models_n \{A\}\}$.
- $\Sigma \models_n^+ \{A\}$ *iff every* n^+*-model of* Σ *is an* n^+*-model of* $\{A\}$ *(on* H*). The set of* n^+*-consequences of* Σ*,* $\mathbf{C}_n^+(\Sigma) = \{A : \Sigma \models_n^+ \{A\}\}$.

For readability we'll write $\Sigma \models_n A$ and $\Sigma \models_n^+ A$ instead. We'll also omit references to the underlying hypergraph. Before we show that $\models_n$ and $\models_n^+$ are genuine consequence relations in the sense of [168] (i.e. reflexive, monotonic, and transitive), we'll first state some obvious facts based on definition (7.4.1):

Fact 7.4.1
For any Σ and A:

1. *every n^+-model of Σ is an n-model of Σ.*
2. *if $\Sigma \models_n A$, then $\Sigma \models_n^+ A$ (equivalently, $\mathbf{C}_n(\Sigma) \subseteq \mathbf{C}_n^+(\Sigma)$)*
3. *$\mathbf{C}_n^+(\Sigma) = \mathbf{C}_n(\Sigma) \cup \mathbf{C}_n(\emptyset)$*

Where B is formula and α is a clause on H, we let

$$B \setminus \alpha = \{\beta \setminus \alpha : \beta \in B\}$$

We can now give a characterisation of $\models_n$ in terms of resolution and Σ-n-colourings.

Theorem 7.4.1
Let Σ be a set of formulae and A a formula on H, then relative to H, $\Sigma \models_n A$ iff for each Σ-n-colouring f there exists an $i \leq n$ such that for each $\alpha \in A$, $\emptyset \in [(\cup(f^{-1}[i])) \setminus \alpha]_H$.

Proof:
($\Rightarrow$) Assume that $\Sigma \models_n A$. Towards a contradiction, let f_0 be a Σ-n-colouring such that for each $i \leq n$ there exists some $\alpha \in A$ with $\emptyset \notin [(\cup(f_0^{-1}[i])) \setminus \alpha]_H$. By theorem (7.3.1), for each $i \leq n$ there exists an $\alpha \in A$ such that $(\cup(f_0^{-1}[i])) \setminus \alpha$ is weakly satisfiable on H. For each $i \leq n$, let Γ_i weakly satisfy $(\cup(f_0^{-1}[i])) \setminus \alpha$. Then each Γ_i also weakly satisfies $\cup(f_0^{-1}[i])$ and hence $\Gamma_1, \ldots, \Gamma_n$ n-satisfy Σ. But for each Γ_i there exists some $\alpha \in A$ such that $\Gamma_i \cap \alpha = \emptyset$, so $\Gamma_1, \ldots, \Gamma_n$ doesn't n-satisfy A.

($\Leftarrow$) Let $\Gamma_1, \ldots, \Gamma_n$ witness that $\Sigma \not\models_n A$. For each $i \leq n$, let $\sigma_i = \{B \in \Sigma : \Gamma \text{ weakly satisfies } B\}$. Without loss of generality, we may assume that σ_i and σ_j, $i \neq j$, are disjoint, and each $\Gamma_i \subseteq \bigcup\bigcup \sigma_i$. Define the Σ-n-colouring f_0 such that for each $A \in \Sigma$, $f_0(A) = i$ iff $A \in \sigma_i$. Towards a contradiction assume that there exists an $i_0 \leq n$ such that for each $\alpha \in A$, $\emptyset \in [(\cup(f_0^{-1}[i_0])) \setminus \alpha]_H$. Then by theorem (7.3.1), for each $\alpha \in A$, $(\cup(f_0^{-1}[i_0])) \setminus \alpha$ is not weakly satisfiable on H. But by the initial assumption Γ_{i_0} must weakly satisfy $\cup(f_0^{-1}[i_0])$ and there must be an $\alpha_0 \in A$ such that $\Gamma_{i_0} \cap \alpha_0 = \emptyset$. Hence Γ_{i_0} weakly satisfies $(\cup(f_0^{-1}[i_0])) \setminus \alpha_0$ on H. ∎

Theorem 7.4.2
$\models_n$ *has the following structural properties:*

R: $A \in \Sigma \Longrightarrow \Sigma \models_n A$

M: $\Sigma \models_n A \Longrightarrow \Sigma, \Delta \models_n A$

T: $\Sigma, A \models_n B$ *and* $\Sigma \models_n A \Longrightarrow \Sigma \models_n B$

Proof:
For [R], we observe that every n-model of Σ is also an n-model of A for every $A \in \Sigma$. For [M], we observe that every n-model of Σ, Δ is also an n-model of Σ. For [T], we observe that every n-model of Σ is an n-model of A and so it is an n-model of Σ, A. Hence, it is also an n-model of B. ∎

Corollary 7.4.1
$\mathbf{C}_n$ *is a closure operator on* Φ.

In light of fact (7.4.1), properties of $\models_n$ and $\mathbf{C}_n$ can be transferred directly to $\models_n^+$ and $\mathbf{C}_n^+$. In terms of the level function λ_H, both $\mathbf{C}_n$ and $\mathbf{C}_n^+$ are λ_H preserving closure operators.

Theorem 7.4.3
For arbitrary but fixed $n \in \mathbb{N}$, *for any* $\Sigma \subseteq \Phi$, $\lambda_H(\Sigma) n \Leftrightarrow \lambda_H(\mathbf{C}_n(\Sigma)) = n$

Proof:
Suppose that $\lambda_H(\Sigma) = n$. Then there must be an n-model, $\Gamma_1, \ldots, \Gamma_n$, of Σ. But any n-model of Σ is an n-model of A for each $A \in \mathbf{C}_n(\Sigma)$, so $\Gamma_1, \ldots, \Gamma_n$ is an n-model of $\mathbf{C}_n(\Sigma)$. Hence, $\lambda_H(\mathbf{C}_n(\Sigma)) \leq n$. But $\lambda_H(\mathbf{C}_n(\Sigma)) \not< n$, lest $\lambda_H(\Sigma) < n$. Thus $\lambda_H(\mathbf{C}_n(\Sigma)) = n$ as required.

Conversely, suppose that $\lambda_H(\mathbf{C}_n(\Sigma)) = n$. By Inclusion, $\Sigma \subseteq \mathbf{C}_n(\Sigma)$ and thus $\lambda_H(\Sigma) \leq n$. Towards a contradiction, suppose that $\lambda_H(\Sigma) = m < n$. Then by the first part of our proof, $\lambda_H(\mathbf{C}_n(\Sigma)) = m < n$. But this contradicts the leastness of n. Hence $\lambda_H(\mathbf{C}_n(\Sigma)) \not< n$, i.e. $\lambda_H(\mathbf{C}_n(\Sigma)) = n$. ∎

Finally it is straightforward to show the lattice theoretic properties of quotient sets formed by equivalence classes of formulae (on H) defined in terms of $\mathbf{C}_n$ and $\mathbf{C}_n^+$.

Definition 7.4.2
Let Σ, $\Sigma' \subseteq \Phi$ *and* $n \in \mathbb{N}$ *be arbitrary but fixed. We define the binary relation* $\equiv$ *over* Φ^2 *by setting* $\Sigma \equiv \Sigma'$ *iff* $\mathbf{C}_n(\Sigma) = \mathbf{C}_n(\Sigma')$. *We let* $[\Sigma] = \{\Sigma' \subseteq \Phi : \Sigma \equiv \Sigma'\}$ *and* $H/\!\equiv\, = \{[\Sigma] : \Sigma \subseteq \Phi\}$. *For any* $[\Sigma], [\Sigma'] \in H/\!\equiv$, *we let* $[\Sigma] \leq [\Sigma']$ *iff* $\mathbf{C}_n(\Sigma') \subseteq \mathbf{C}_n(\Sigma)$

Theorem 7.4.4

Let $H/\equiv$ *and* $\leq$ *be as defined in definition (7.4.2). Then* $\mathfrak{L} = \langle H/\equiv, \leq \rangle$ *is a complete distributive lattice with a minimum:* $[\Sigma] = 0$ *iff* $\mathbf{C}_n(\Sigma) = \mathbf{C}_n(\Phi)$*. If* $\mathbf{C}_n$ *is replaced with* $\mathbf{C}_n^+$ *throughout in definition (7.4.2), then* $\mathfrak{L}$ *is a complete distributive lattice with both a minimum and maximum, in particular:* $[\Sigma] = 1$ *iff* $\mathbf{C}_n(\Sigma) = \mathbf{C}_n(\emptyset)$

Proof:

In light of fact (7.4.1), we only need to consider $\mathbf{C}_n$. It is straightforward to verify that $\equiv$ is an equivalence relation on Φ and thus every element of $H/\equiv$ is an equivalence class modulo $\equiv$. Moreover, $\leq$ is reflexive, antisymmetric, and transitive. Thus $\leq$ is a partial ordering on $H/\equiv$.

Let I be an index set of arbitrary cardinality. Let $\{[\Sigma_i] : i \in I\} \subseteq H/\equiv$. We'll show that $[\bigcup_{i \in I} \mathbf{C}_n(\Sigma_i)]$ and $[\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)]$ are, respectively, the greatest lower bound and the least upper bound of $\{[\Sigma_i] : i \in I\}$:

(1) $[\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)]$ is an upper bound:

$$\begin{aligned}
\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i) \subseteq \mathbf{C}_n(\Sigma_j) \text{ for each } j \in I \quad &\Longrightarrow \quad \mathbf{C}_n(\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)) \subseteq \mathbf{C}_n(\mathbf{C}_n(\Sigma_j)) \\
& \quad \text{for each } j \in I \\
&\Longrightarrow \quad \mathbf{C}_n(\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)) \subseteq \mathbf{C}_n(\Sigma_j) \\
& \quad \text{for each } j \in I \\
&\Longrightarrow \quad [\Sigma_j] \leq [\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)] \\
& \quad \text{for each } j \in I
\end{aligned}$$

(2) $[\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)]$ is the least upper bound:

$$\begin{aligned}
[\Sigma_i] \leq [\Delta] \text{ for each } i \in I \quad &\Longrightarrow \quad \mathbf{C}_n(\Delta) \subseteq \mathbf{C}_n(\Sigma_i) \text{ for each } i \in I \\
&\Longrightarrow \quad \mathbf{C}_n(\Delta) \subseteq \bigcap_{i \in I} \mathbf{C}_n(\Sigma_i) \\
&\Longrightarrow \quad [\bigcap_{i \in I} \mathbf{C}_n(\Sigma_i)] \leq [\Delta]
\end{aligned}$$

(3) $[\bigcup_{i \in I} \mathbf{C}_n(\Sigma_i)]$ is a lower bound: similar to (1).

(4) $[\bigcup_{i \in I} \mathbf{C}_n(\Sigma_i)]$ is the greatest lower bound: similar to (2).

$\mathfrak{L}$ is thus a complete lattice. The distributivity of $\mathfrak{L}$ follows from the fact that $\cap$ is distributive over $\cup$ and vice versa. To verify that $[\Phi]$ is the minimum, we observe that $\mathbf{C}_n(\Sigma) \subseteq \mathbf{C}_n(\Phi)$ and $[\Phi] \leq [\Sigma]$ for any Σ. For the case of $\mathbf{C}_n^+$, since $\mathbf{C}_n^+(\emptyset) \subseteq \mathbf{C}_n^+(\Sigma)$ it follows that $[\Sigma] \leq [\emptyset]$ for any Σ. Hence $[\emptyset]$ is the maximum. ∎

Question 7.4.1
Are $\mathbf{C}_n$ and $\mathbf{C}_n^+$ algebraic closure operators, i.e. are $\models_n$ and $\models_n^+$ compact?

7.5 BPI and Complexity Theory

In section (7.2) we demonstrate that the compactness statement for n-satisfiability on hypergraphs is equivalent to BPI in ZF set theory without the axiom of choice. In section (7.3), we note further that the corresponding decision problem for n-satisfiability on compact hypergraphs is **NP**-complete. Our investigation is partly motivated by a conjecture from [50; 52]. Cowen notices that methods for proving certain decision problems are **NP**-complete have also been used in showing that certain compactness theorems are equivalent to BPI in ZF. More specifically, let R be a compactness statement that says of a set S and a property P that if every finite *subobject* of an object in S has P, then the object has P. Moreover, we assume that R is not equivalent in ZF to the statement that every object in S has P. This additional assumption is required to eliminate certain bogus compactness statements. If R is a compactness statement in the above sense, R^* will denote the corresponding decision problem which asks of a finite object in S, does it have the property P; $R < \mathrm{BPI}$ will denote that R is weaker than BPI in ZF, and $R \Leftrightarrow \mathrm{BPI}$ will denote that R is equivalent to BPI in ZF. In [50] Cowen gives various examples for R and R^*. But all of Cowen's examples fall into 3 types: (1) $R \Leftrightarrow \mathrm{BPI}$ and R^* is **NP**-complete, (2) $R < \mathrm{BPI}$ and R^* is polynomial, and (3) $R < \mathrm{BPI}$ and R^* is **NP**-complete.

	$R \Leftrightarrow \mathrm{BPI}$	$R < \mathrm{BPI}$
R^* is polynomial	?	+
R^* is **NP**-complete	+	+

Figure 7.2: Known cases of R and R^* in relation to BPI

But there is no known example of R where $R \Leftrightarrow \mathrm{BPI}$ while R^* is polynomial. Cowen makes the following conjecture:

Conjecture 7.5.1
(Cowen) If R^* is polynomial, then $R < \mathrm{BPI}$.

Cowen's conjecture implies, in particular, that $\mathbf{P} \neq \mathbf{NP}$ since letting R be the compactness statement for n-satisfiability on hypergraphs gives $\mathsf{R} \Leftrightarrow \mathsf{BPI}$. But R^* would be polynomial if $\mathbf{P} = \mathbf{NP}$. Hence any proof of Cowen's conjecture would be a *de facto* proof that $\mathbf{P} \neq \mathbf{NP}$.

The conjecture of Cowen also opens a new line of inquiry: of a particular polynomial R^*, we can ask whether the corresponding $\mathsf{R} < \mathsf{BPI}$ holds. To take one particular example it is known that 2-**SAT**, i.e. satisfiability of clauses with at most 2 literals, is polynomial; the corresponding compactness statement for 2-**SAT** has been shown by Wojtylak in [181] to be weaker then BPI.

Question 7.5.1
Let R *be the statement that a collection of finite sets has a system of distinct representatives iff every finite subcollection has a system of distinct representatives. It is known that the corresponding* R^* *is polynomial. But is* $\mathsf{R} < \mathsf{BPI}$*?*

Chapter 8

Conclusion

The study of logic usually begins with one of two approaches. According to what Priest [139] calls the *canonical* approach, the aim of logic is to establish a standard for evaluating arguments – a standard by which we judge whether a conclusion can be legitimately inferred from a body of assumptions. The legitimacy of an inference turns on the notion of a consequence relation which can be defined proof theoretically in terms of deduction or semantically in terms of class containment of models. Legitimate or valid inferences are those that are sanctioned by our consequence relation specified in standard proof theoretic or model theoretic ways. The usual completeness theorem of a logic is in turn an assurance that the proof theory and the semantics capture one and the same consequence relation.

According to the *representational* approach however, logic is understood as the study of the relationship between a formal language and its associated domains. It addresses issues concerning *how* to express and *what* can be expressed in a formal language. Although the two approaches have different aims, they are clearly related. Amongst the sort of things we want a formal language to be able to express are *declarative* sentences (in contrast to *imperative* sentences, e.g. `goto s` in some programming languages). The content of these declarative sentences is fixed by their *truth conditions* which in turn inform us that entities or states in the domain are one way but not another. The availability of sentences bearing truth conditions allows us to be in the business of reasoning and inference again. If a body of declarative sentences truthfully represents the domain, we can infer further truthful sentences about the domain.

Given these two approaches to logic, it is not surprising that paraconsistent logics are typically motivated in one of two ways. According to the epistemic account, declarative sentences in a formal language can be used to represent states of a domain, in particular they can be used to represent states of the actual world. We may think of these representations as logical descriptions with empirical contents. Although we are the masters of our own language, infallibility is not a given. We make mistakes and some of them turn up as inconsistencies in our data and theories. As Wheeler pointed out in [180], some of these mistakes, e.g. measurement errors, are so fundamental to the way we interact with the world that any attempt to eliminate them is a practical impossibility. Our scientific theories and thus scientific reasoning must face up to the force of inconsistencies. According to this view certain inconsistencies are just misrepresentations. Since logic is about consequence, it is the logician's business to sort out what can be deduced from these misrepresentations. Classical logic is of no help here since it does not distinguish between different sorts of mistakes and hence all sorts of mistakes can be inferred. Adopting alternative logics is one way we can continue to draw inferences under the threat of erroneous data.

According to the ontological account however, not all inconsistent descriptions are infected with errors or misrepresentations. Instead an alternative hypothesis is that certain inconsistent descriptions are just *correct* descriptions of entities or states that are inconsistent in and of themselves. As Priest would say, some inconsistent information or theories are *true* ([140]). According to the dialetheic thesis the recurrence of certain paradoxical statements is not to be explained away in terms of mistakes on our part or defects in our language. Instead *the best explanation* of the persistence of these paradoxes is that they truly describe an ontology populated with inconsistent entities. The implication of the ontological or dialetheic account is that just as we need alternative logics to reason with potentially erroneous data or theories, we also need alternative logics to reason with inconsistent entities and states. Once again, classical logic is of no help since it provides no provision to deal with inconsistent entities or states.

While dialetheism, the hypothesis that there are true inconsistent theories, is a contentious claim, we have neither defended nor criticised dialetheism in this thesis. We have taken it for granted that the epistemic account is a plausible motivation for paraconsistency and investigated a variety of inconsistency-tolerant reasoning strategies. Nonetheless we are in agreement with the dialetheist that not every case of inconsistent description is a case of misrepresentation or error. We have pointed out that there are situations where the source of inconsistencies is neither rooted in error nor in inconsistent ontologies. In drawing out these cases, our emphasis is on the practical use of a logic as a formalism for *representation*. We have so far steered clear of any undue ontological commitment to endorse inconsistent entities or epistemic pressure to eradicate inconsistencies. Our *modus operandi* is preservationalism – we look for useful properties of inconsistent descriptions and find inference strategies that preserve these properties.

In this thesis we have highlighted the fact that logic is as much about *representation* as it is about *consequence*. In viewing logic as a language for modelling practical and abstract problems, the emphasis is on the discriminatory power of our language. The main issue in this thesis is not merely fault tolerant deduction per se or the reality of an inconsistent ontology. Rather, it is the analysis of the structure and the underlying combinatorial properties of our logical representation which in turn inform us about the nature of the situation or problem under consideration. Of course to provide such an analysis, our logical description must capture the salient features of the problem *at some appropriate level of abstraction*. But this is very much a question about the representational efficacy of the formal language *and* the representational fit between the formal language and the problem domain. In saying this, we do not intend to suggest that representational issues have *nothing* to do with deduction. Far from it, deduction is related to meaning. As is well known, the meaning of a logical connective can be specified by the use of introduction and elimination rules. So deduction can be used to ground and fix the meaning of a logical representation. But note that this way of bringing deduction back into the picture requires no tacit assumption about epistemic error or inconsistent ontology. We maintain that there are cases in which a problem domain is best modelled by logically inconsistent descriptions involving neither epistemic error nor ontological assertion. Of course the representational fit between a logical description and a problem domain must be evaluated in the context of a machinery for specifying the meaning of the description. In this work we have not committed to any particular way to accomplish the task. There is no harm in being a methodological pluralist. Whether one opts for a model theoretic or proof theoretic machinery, incompleteness and unsoundness are genuine possibilities. There is no guarantee that a model theoretic specification must have a corresponding proof theoretic specification or vice versa.

The usefulness of paraconsistent logics as a way to ground the meaning of formal languages is perhaps analogous to the usefulness of a scientific instrument. Ancient astronomers carved out the constellations with their bare eyes, charting the night sky into distinct heavenly bodies and regions. They did what they could given what was available at the time. Astronomers in the Renaissance were bestowed with the gift of the telescope. They could now chart the night sky with finer precision and distinctions that were not seen by ancient astronomers. Modern astronomers go one step further by tapping into the unexplored territory of radio frequencies. Formal languages are the symbolic constellations for the modern logician. The history of modern logic too is punctuated with remarkable changes in the discriminatory and expressive power of logics. Propositional logic delivers the calculus of *propositions*. But to formalise 'Every natural number has a successor', we had to await for the advent of quantification theory. To formalise 'A relation is well-ordered if every non-empty subset has a least element', we have to go second-order or employ set theory of some form. [1] At each turn of this refinement, more can be said and more can be discerned. But at the same time it is also surprisingly conservative. In classical logics, sentences are sorted into three distinct classes – those that are tautologous (true in all models), inconsistent (false in all models), and contingent (true in some models and false in some models). [2] Even in a simple classical propositional logic with countably many propositional atoms, there are at least countably many distinct non-equivalent classes of contingent sentences – one for each distinct atom. But oddly, there can only be *one* equivalence class of tautologous sentences and *one* equivalence class of inconsistent sentences. Within the classical scheme a contradiction, $(p \wedge \neg p)$, is indistinguishable from a denial of the excluded middle, $\neg(q \vee \neg q)$ – no classical model, no classical proof will separate them. But note that we do distinguish these sentences meta-logically. They are not merely distinct syntactic tokens of distinct types – while one can be used to assert a contradiction the other can be used to reject the law of excluded middle. The distinction also carries a certain semantic weight.

Our take home message then is this: formal languages which express inconsistencies are rich in structure and expressive power. Our complaint against classical consequence and classical semantics is that they do not make room for the discrimination of different types of inconsistencies within a formal language. Under the classical scheme, all inconsistencies are proof-theoretically and model-theoretically *equivalent*. But recall that the study of formal languages, standard first order model theory in particular, is very much concerned with the discriminatory power of languages and models. To make distinctions, we must be able to partition a language into distinct equivalence classes. In fact, in a very general sense all formal languages are concerned with equivalence classes – namely classes that are organised under the 'sameness-in-

[1] We dare not say 'the set theory' here. As we all know ZF is to be distinguished from ZFC (with choice axiom), from ZFA (with anti-foundation axiom) from NBG (NBG for von Neumman, Bernays and Godel, not to be confused with the epithet 'No Bloody Good'), from (Quine's) NF. Set theories, like logics, come in many varieties.

[2] Given soundness and completeness of classical logics, the reader may use the appropriate proof-theoretic substitutes for 'true in all models' etc.

meaning' relation. Indeed this is one of the main goals (and advantages) of the study of formal languages – given any two expressions in a formal language we want to provide a systematic and rigorous method to determine if the two have the same meaning. Logic provides a paradigmatic method of doing this – in fact it provides two methods, one via proofs the other via models.[3] So the inability of the classical scheme to discern different inconsistencies is a failing on its part to do its job.

Viewing the matter in this light gives us the satisfaction of putting a positive spin on paraconsistency and turning the tables on the classicist. It is often said that paraconsistent logics are simply too weak to do any real work – they give up too many classically acceptable rules of inference. Our rejoinder is that sometimes weakness is also a strength. Recall that in the study of modalities, any finite sequence of $\neg$, $\Box$, and $\Diamond$, strong modal logics such as S4 and S5 have finitely many modalities. More precisely, S4 has 14 distinct modalities while S5 has only 6. So in terms of the discriminatory power of these logics, we can only express 14 and 6 distinct types of modal statements. These logics are strong, but they don't necessarily give us greater distinction. The comparison between classical and paraconsistent logics is analogous to the comparison between strong and weak modal logics. Strength in deducibility is not tantamount to strength in discriminatory power. Paraconsistent logics and semantics are not merely non-explosive, they also allow us to preserve important distinctions. Not all inconsistencies are equal and they should not be. Paraconsistent logics are endowed with the power to discriminate between different inconsistencies. This, we maintain, is another way to 'go beyond consistency'.

[3]We take this to be at least a necessary condition for such a semantic specification. However, it is debatable whether it is also sufficient. Some may insist that the equivalence relation induced by the underlying logic must also be a *congruence* relation. This amounts to the requirement that intersubstitutivity of provable equivalents preserves equivalence. As is well known many paraconsistent logics e.g. Priest's LP and da Costa's C-systems, do not have such a property. We do not wish to settle the issue here. But we do think that it is a research direction worth further investigation.

Appendix A

Dunn's Ambi-Valuation Semantics

Definition A.0.1

Let A and B be truth functional formulae (i.e. zero degree formulae). A tautologically entails B, $A \rightarrow B$, iff there is a disjunctive normal form (DNF) of $A = A_1 \vee \ldots \vee A_n$ and there is a conjunctive normal form (CNF) of $B = B_1 \wedge \ldots \wedge B_m$ such that for each $i \leq n$ and $j \leq m$, A_i is a term and B_j is a clause and A_i and B_j have a common literal.

In Belnap [7], it is shown that the set of tautological entailments is precisely the first degree fragment of the relevant logics E and R.

Definition A.0.2

A relevant assignment v is a function $v : At \longrightarrow \wp\{1,0\}$. A relevant assignment v is extended uniquely to $\bar{v}$ over Φ by the following recursion:

1. $\bar{v}(p_i) = v(p_i)$ *for any* $p_i \in At$
2. $1 \in \bar{v}(\neg A) \;\Leftrightarrow\; 0 \in \bar{v}(A)$
   $0 \in \bar{v}(\neg A) \;\Leftrightarrow\; 1 \in \bar{v}(A)$
3. $1 \in \bar{v}(A \wedge B) \;\Leftrightarrow\; 1 \in \bar{v}(A)$ *and* $1 \in \bar{v}(B)$
   $0 \in \bar{v}(A \wedge B) \;\Leftrightarrow\; 0 \in \bar{v}(A)$ *or* $0 \in \bar{v}(B)$
4. $1 \in \bar{v}(A \vee B) \;\Leftrightarrow\; 1 \in \bar{v}(A)$ *or* $1 \in \bar{v}(B)$
   $0 \in \bar{v}(A \vee B) \;\Leftrightarrow\; 0 \in \bar{v}(A)$ *and* $0 \in \bar{v}(B)$

A relevant valuation $\bar{v}$ is a model of A iff $1 \in \bar{v}(A)$. We write $A \models_R B$, A relevantly entails B, iff every relevant model of A is a relevant model of B.

Alternatively, we can define $\models_R$ using the four-valued matrices in figure (A.1) instead. We use T for $\{1\}$, F for $\{0\}$ and B for $\{1,0\}$ and N for $\emptyset$. Taking T and B as the designated values, we can define $A \models_R B$ iff $v(A) \in \{T, B\} \implies v(B) \in \{T, B\}$. The lattice **4** interpreted as subsets of $\{1,0\}$ is given in figure (A.2).

¬	T	F	B	N
	F	T	B	N

∨	T	F	B	N
T	T	T	T	T
F	T	F	B	N
B	T	B	B	T
N	T	N	T	N

∧	T	F	B	N
T	T	F	B	N
F	F	F	F	F
B	B	F	B	F
N	N	F	F	N

Figure A.1: 4-valued matrices for FDE

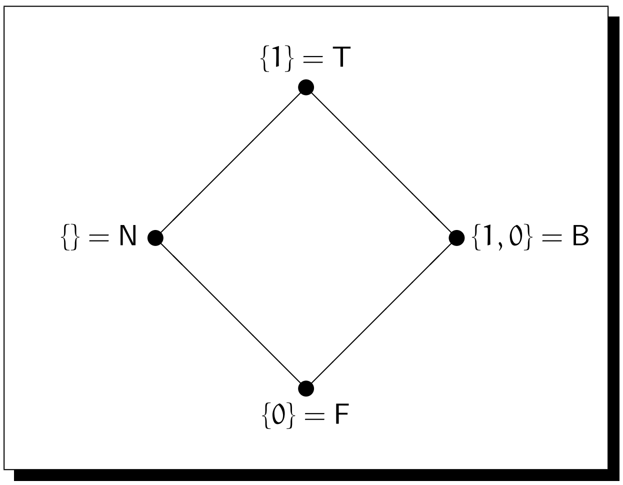

Figure A.2: The lattice **4** interpreted as subsets of $\{1, 0\}$.

Theorem A.0.1
(Belnap[7], Dunn [65]) $A \rightarrow B \;\Leftrightarrow\; A \models_R B$.

Alternatively, FDE can also be characterised as entailment between clauses (see Hanson [83] and Levesque [119]). For a clause or a term A we use $\mathrm{lit}(A)$ to denote the set of literals occurring in A.

Definition A.0.3
Let A *and* B *be truth functional formulae. Then* A *clausally entails* B, $A \rightarrow_c B$ *iff there is a* CNF *of* A, $A_1 \wedge \ldots \wedge A_n$, *and a* CNF *of* B, $B_1 \wedge \ldots \wedge B_m$, *such that for every* $i \leq m$ *there is some* $j \leq n$ *with* $\mathrm{lit}(A_j) \subseteq \mathrm{lit}(B_i)$.

We note that if A and B are already in CNF, then it takes only $\mathcal{O}(|A| \cdot |B|)$ time to determine whether $A \rightarrow_c B$. Indeed this is the main reason why Levesque [119] finds clausal entailment an attractive model of quick surface reasoning of an agent.

Proposition A.0.1

The set of tautological entailments is exactly the set of clausal entailments, i.e. for any truth functional A *and* B, $A \rightarrow B \;\Leftrightarrow\; A \rightarrow_c B$

Proof:

($\Rightarrow$) Assume that $A \rightarrow B$ and let $A' = A_1 \vee \ldots \vee A_n$ and $B' = B_1 \wedge \ldots \wedge B_m$ be the witness. We define A^* as follows:

$$A^* = \bigwedge_{l_{A_i} \in lit(A_i)}^{i \leq n} (l_{A_1} \vee \ldots \vee l_{A_n})$$

i.e. A^* is a formula in CNF, $C_1 \wedge \ldots \wedge C_k$, where each clause C_j is composed of literals from distinct terms of A'. It is easy to verify that for every $j \leq m$ there is some $i \leq k$ such that $lit(C_i) \subseteq lit(B_j)$ (since all possible combinations of literals from distinct terms of A' are represented by clauses of A^*). Hence, $A \rightarrow_c B$ as required.

($\Leftarrow$) Assume that $A \rightarrow_c B$ and let $A'' = A_1 \wedge \ldots \wedge A_n$ and $B'' = B_1 \wedge \ldots \wedge B_m$ be the witnesses. We define A^{**} similarly as follows:

$$A^{**} = \bigvee_{l_{A_i} \in lit(A_i)}^{i \leq n} (l_{A_1} \wedge \ldots \wedge l_{A_m})$$

It is straightforward to verify that every term of A^{**} and every clause of B'' have a common literal. Hence $A \rightarrow B$ as required. ∎

Appendix B

The Pair Extension Lemma in Analytic Implicational Logics

The role of the Pair Extension lemma in the completeness proof for relevant logics (see [8]) is analogous to the role of Lindenbaum's lemma in Henkin completeness proofs for classical and modal logics. The Pair Extension lemma is in fact a very natural generalisation of Lindenbaum's lemma. In the words of Dunn, the Pair Extension (or Belnap's) lemma 'symmetrizes the usual Henkin construction of 1st order classical logic' ([66] p. 160).[1] Lindenbaum's lemma says that an L-consistent set of formulae can always be extended consistently to maximality. The Pair Extension lemma says that an L-*exclusive* pair of sets of formulae can always be extended, L-exclusively, to a pair of sets that is also L-*exhaustive*. The proofs of both of these lemmata require constructions, from the original set(s), that can preserve either L-consistency or L-exclusivity. In these constructions, a certain lattice property of disjunction is assumed. In particular the axiom of addition, $A \rightarrow A \vee B$, must be a theorem of the logic (see p. 121 [8]). In this note we'll show that for certain logics without $A \rightarrow A \vee B$ as a theorem we can still prove the Pair Extension lemma.

An implication, $A \rightarrow B$, is said to be *analytic* if all of B's sentential variables are included in A's sentential variables. An analytic implicational logic is one in which all implicational theorems are analytic. The first axiomatisation of such a logic is given by Parry [135]. In [64], Dunn modified Parry's system by *demodalising* the system. Urquhart [178] then studied a modal extension of Dunn's system. The completeness of Parry's original system is finally proven by Fine in [69] thereby answering a question of Gödel. In more recent years, certain *paraconsistent* versions of analytic implicational logics have been studied by Deutsch [59; 60; 61; 62] and Sylvan (formerly Routley)

[1] While Dunn attributes the lemma to Belnap, Gabbay [71] gave an independent proof of the analogue for first order intuitionistic logic with constant domain.

[176]. The main interest in these newer logics is that they combine features of both analytic and relevant implication.

We say that an implicational logic L is *pair extension acceptable* if it satisfies the following conditions:

Modus ponens $\vdash_L A \rightarrow B$ and $\vdash_L A \Longrightarrow \vdash_L B$

Reflexivity $\vdash_L A \rightarrow A$

Transitivity $\vdash_L A \rightarrow B$ and $\vdash_L B \rightarrow C \Longrightarrow \vdash_L A \rightarrow C$

Conjunction

(a) $\vdash_L A \wedge B \rightarrow A \qquad \vdash_L A \wedge B \rightarrow B$,

(b) $\vdash_L A \rightarrow B$ and $\vdash_L A \rightarrow C \Longrightarrow \vdash_L A \rightarrow B \wedge C$

(c) $\vdash_L A$ and $\vdash_L B \Longrightarrow \vdash_L A \wedge B$

Disjunction $\vdash_L A \rightarrow B$ and $\vdash_L C \rightarrow D \Longrightarrow \vdash_L A \vee C \rightarrow B \vee D$

Distribution $\vdash_L A \wedge (B \vee C) \rightarrow (A \wedge B) \vee C$

We note that the only difference between this definition and the definition of up-down acceptability given in [8] (p.121) is the property of disjunction. Here we require neither $\vdash_L A \rightarrow A \vee B$ nor $\vdash_L B \rightarrow A \vee B$. This is in line with the main drift of analytic implication. For up-down acceptable logics, disjunction has the following properties:

($\vee$R) $\vdash_L A \rightarrow A \vee B \qquad \vdash_L B \rightarrow A \vee B$

($\vee$L) $\vdash_L A \rightarrow C$ and $\vdash_L B \rightarrow C \Longrightarrow \vdash_L A \vee B \rightarrow C$

Proposition B.0.2

Any implicational logic L *that satisfies (*$\vee$R*), (*$\vee$L*) and* ***Transitivity*** *above also satisfies* ***Disjunction****.*

Proof:

(1)	$\vdash_L A \rightarrow B$	(Assumption)
(2)	$\vdash_L C \rightarrow D$	(Assumption)
(3)	$\vdash_L B \rightarrow B \vee D$	($\vee$R)
(4)	$\vdash_L D \rightarrow B \vee D$	($\vee$R)
(5)	$\vdash_L A \rightarrow B \vee D$	(1, 3 Transitivity)
(6)	$\vdash_L C \rightarrow B \vee D$	(2, 4 Transitivity)
(7)	$\vdash_L A \vee C \rightarrow B \vee D$	(5, 6 $\vee$L)

∎

We should also note that, with the exception of **Modus ponens**, all stated conditions are *variable preserving*. So for instance if $A \to B$ and $C \to D$ are both theorems of an analytic implicational logic, then $A \vee C \to B \vee D$ will also pass the variable containment requirement.

The following proposition shows that being an analytic implicational logic is compatible with being a pair extension acceptable logic.

Proposition B.0.3
There are analytic implicational logics that are pair extension acceptable.

Proof:
The logic S' defined by Deutsch in [60] is analytic, moreover modus ponens is admissible in S'. It is straightforward to verify that all the remaining rules of pair extension acceptable logics are validity preserving in the corresponding Kripke semantics defined for S'. Hence by the completeness result of Deutsch all of the rules are derivable in S'. ∎

We now introduce some key definitions. Let $\Gamma,\ \Gamma',\ \Delta,\ \Delta'$ be sets of formulae and let L be a logic. Then we say that an ordered pair $\langle \Gamma', \Delta' \rangle$ extends $\langle \Gamma, \Delta \rangle$ iff $\Gamma \subseteq \Gamma'$ and $\Delta \subseteq \Delta'$. A pair $\langle \Gamma, \Delta \rangle$ is said to be L-exclusive if for no $A_1, \ldots, A_n \in \Gamma$ and $B_1, \ldots, B_m \in \Delta$ do we have $\vdash_L A_1 \wedge \ldots \wedge A_n \to B_1 \vee \ldots \vee B_m$. A pair $\langle \Gamma, \Delta \rangle$ is exhaustive if $\Gamma \cup \Delta$ is the entire language Φ. We are now in a position to prove the following key lemma:

Lemma B.0.1
Let L be a pair extension acceptable logic. If $\langle \Gamma, \Delta \rangle$ is a L-exclusive pair, then for any formula C either $\langle \Gamma \cup \{C\}, \Delta \rangle$ is L-exclusive or $\langle \Gamma, \Delta \cup \{C\} \rangle$ is L-exclusive.

Proof:
Let $\langle \Gamma, \Delta \rangle$ be a L-exclusive pair. Towards a contradiction we assume that neither $\langle \Gamma \cup \{C\}, \Delta \rangle$ nor $\langle \Gamma, \Delta \cup \{C\} \rangle$ is L-exclusive. Then there must be some A, A', B and B' such that

1. $A = A_1 \wedge \ldots \wedge A_i$ and $A' = A'_1 \wedge \ldots \wedge A'_j$ where $A_1 \ldots A_i, A'_1 \ldots A'_j \in \Gamma$;

2. $B = B_1 \wedge \ldots \wedge B_m$ and $B' = B'_1 \wedge \ldots \wedge B'_n$ where $B_1 \ldots B_m, B'_1 \ldots B'_n \in \Delta$;

3. (a) $\vdash_L A \to C \vee B$ and (b) $\vdash_L A' \wedge C \to B'$

The following proof suffices to show that (3a) and (3b) implies that $A \wedge A' \to B' \vee B$ is a L-theorem which contradicts the L-exclusivity of $\langle \Gamma, \Delta \rangle$.

(1)	$\vdash_L A \to C \vee B$	(Assumption 3a)
(2)	$\vdash_L A' \wedge C \to B'$	(Assumption 3b)
(3)	$\vdash_L A \wedge A' \to A$	(Conjunction a)
(4)	$\vdash_L A \wedge A' \to C \vee B$	(1, 3 by Transitivity)
(5)	$\vdash_L A \wedge A' \to A'$	(Conjunction a)
(6)	$\vdash_L A \wedge A' \to A' \wedge (C \vee B)$	(4, 5 Conjunction b)
(7)	$\vdash_L A' \wedge (C \vee B) \to (A' \wedge C) \vee B$	(Distribution)
(8)	$\vdash_L A \wedge A' \to (A' \wedge C) \vee B$	(6, 7 Transitivity)
(9)	$\vdash_L B \to B$	(Reflexivity)
(10)	$\vdash_L (A' \wedge C) \vee B \to B' \vee B$	(2, 9 Disjunction)
(11)	$\vdash_L A \wedge A' \to B' \vee B$	(9, 10 Transitivity)

∎

The key to our proof of the lemma is in line (10) where we appeal to a weaker condition on disjunction. We should also note that, with suitable modification of our definitions, we can completely recast our proof in terms of the consequence relation of L.

We can now officially record the Pair Extension Lemma. The proof is standard, but we include it for completeness sake.

Theorem B.0.2
Pair Extension Lemma: Let L *be a pair extension acceptable logic and* $\langle \Gamma, \Delta \rangle$ *be a* L*-exclusive pair. Then* $\langle \Gamma, \Delta \rangle$ *can be extended to a* L*-exclusive and exhaustive pair* $\langle \Gamma', \Delta' \rangle$.

Proof:
Without loss of generality we may assume that Φ is countable. As usual we give a fixed enumeration of formulae, $A_1, A_2, A_3 \ldots$. We then define a sequence of pairs $\langle \Gamma_0, \Delta_0 \rangle \ldots \langle \Gamma_n, \Delta_n \rangle \ldots$ where $\langle \Gamma_0, \Delta_0 \rangle = \langle \Gamma, \Delta \rangle$, and given $\langle \Gamma_n, \Delta_n \rangle$ we define

$$\langle \Gamma_{n+1}, \Delta_{n+1} \rangle = \begin{cases} \langle \Gamma_n \cup \{A_n\}, \Delta_n \rangle & \text{if } \langle \Gamma_n \cup \{A_n\}, \Delta_n \rangle \text{ is L-exclusive} \\ \\ \langle \Gamma_n, \Delta_n \cup \{A_n\} \rangle & \text{otherwise} \end{cases}$$

It is straightforward to verify that $\langle \Gamma', \Delta' \rangle = \langle \bigcup_{n \in \omega} \Gamma_n, \bigcup_{n \in \omega} \Delta_n \rangle$ is a L-exclusive extension of $\langle \Gamma, \Delta \rangle$. A simple induction on n and our previous lemma guarantees that $\langle \Gamma', \Delta' \rangle$ is L-exclusive. Clearly it is also exhaustive.

An alternative proof can be given using Zorn's lemma without the assumption that Φ is countable. Let $\mathcal{C}$ be the set of L-exclusive pairs $\langle \Gamma', \Delta' \rangle$ such that $\Gamma \subseteq \Gamma'$ and $\Delta \subseteq \Delta'$. We note that $\mathcal{C}$ is not empty since $\langle \Gamma, \Delta \rangle \in \mathcal{C}$. Then partially order $\mathcal{C}$ by $\leq$ where $\langle \Gamma_i, \Delta_i \rangle \leq \langle \Gamma_j, \Delta_j \rangle$ just in case $\Gamma_i \subseteq \Gamma_j$ and $\Delta_i \subseteq \Delta_j$. Clearly the union of any $\leq$-chain is an element in $\mathcal{C}$. Hence, every $\leq$-chain has an upper bound and thus by Zorn's lemma, there is a $\leq$-maximal element in $\mathcal{C}$. It is straightforward to verify that this maximal element is the required extension of $\langle \Gamma, \Delta \rangle$. ∎

Appendix C

List of Publications

'Modal (Logic) Paraconsistency' with Philippe Besnard,
Proceedings of the Seventh European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty, July 2-5, 2003, Aalborg, Denmark; page 540–511, Lecture Notes in Artificial Intelligence 2711, Springer-Verlag

'Reasoning and Modeling: Two Views of Inconsistency Handling', (under review)
Proceedings of the Third World Congress of Paraconsistency, IRIT, Toulouse, France 28–31 July 2003

'Paraconsistent Reasoning as an Analytic Tool', with Philippe Besnard
The Proceedings of the International Conference on Formal and Applied Practical Reasoning, Imperial College, London, 18 – 20 Sept. 2000, *Logic Journal of the Interest Group in Pure and Applied Logics*, volume 9, no. 2, page 233–246, 2001

'Inconsistency and Preservation',
PRICAI 2000 Topics in Artificial Intelligence, 6th Pacific Rim International Conference on Artificial Intelligence, Melbourne, August/September 2000 Proceedings; page 50–60, Lecture Notes in Artificial Intelligence 1886, Springer-Verlag

'From Weak Satisfiability to n-Satisfiability on Hypergraphs',
The Proceedings of the 12th European Summer School in Logic, Language and Information, Student Session, University of Birmingham, 6–18 August 2000, page 275-285 (CD ROM)

Bibliography

[1] E. W. Adams. Transmissible Improbabilities and Marginal Essentialness of Premises in Inferences Involving Indicative Conditionals. *Journal of Philosophical Logic*, 10:149–177, 1981.

[2] E. W. Adams. On the Logic of High Probability. *Journal of Philosophical Logic*, 155:255–279, 1986.

[3] E. W. Adams. Four Probability-Preserving Properties of Inferences. *Journal of Philosophical Logic*, 25:1–24, 1996.

[4] E. W. Adams. *A Primer of Probability Logic*. CSLI, 1998.

[5] E. W. Adams and H. P. Levine. On the Uncertainties Transmitted from Premises to Conclusions in Deductive Inferences. *Synthese*, 30:429–460, 1975.

[6] J. Aisbett and G. Gibbon. A Practical Measure of the Information in a Logic Theory. *Journal of Experiment and Theoretical Artificial Intelligence*, 11:201–217, 1999.

[7] A. E. Anderson and N. D. Belnap. *Entailment: the Logic of Relevance and Necessity Vol 1*. Princeton University Press, 1975.

[8] A. E. Anderson, N. D. Belnap, and J. M. Dunn. *Entailment: the Logic of Relevance and Necessity Vol 2*. Princeton University Press, 1992.

[9] P. J. Apostoli. Modal Aggregation and the Theory of Paraconsistent Filters. *Mathematical Logic Quarterly*, 42:175–190, 1996.

[10] P. J. Apostoli. On the Completeness of First Degree Weakly Aggregative Modal Logics. *Journal of Philosophical Logic*, 26:169–180, 1997.

[11] P. J. Apostoli and B. Brown. A Solution to the Completeness Problem for Weakly Aggregative Modal Logic. *Journal of Symbolic Logic*, 60 (3):832–842, 1995.

[12] O. Arieli and A. Avron. Reasoning with Logical Bilattices. *Journal of Logic, Language, and Information*, 5:25–63, 1996.

[13] O. Arieli and M. Denecker. Modelling Paraconsistent Reasoning by Classical Logic. In T. Eiter and K-D. Schewe, editors, *Proceedings of the Second International Sympo-*

sium on Foundations of Information and Knowledge Systems (FoIKS), Lecture Notes in Computer Science 2284, pages 1–14. Springer Verlag, 2002.

[14] O. Arieli and M. Denecker. Reducing Preferential Paraconsistent Reasoning to Classical Entailment. *Journal of Logic and Computation*, 13 (4):557–580, 2003.

[15] A. Avron and B. Konikowska. Decomposition Proof Systems for Gödel-Dummett Logics. To appear in Studia Logica, 2000.

[16] J. Barwise. *The Situation in Logic*. CSLI, 1989.

[17] J. Barwise. Information and Impossibilities. *Notre Dame Journal of Formal Logic*, 38 (4):488–515, 1997.

[18] R. J. Bayardo Jr. and J. D. Pehoushek. Counting Models Using Connected Components. In *Proceeding of the National Conference on Artificial Intelligence (AAAI-2000)*, 2000.

[19] J. L. Bell. Polymodal Lattices and Polymodal Logic. *Mathematical Logic Quarterly*, 42:219–233, 1996.

[20] J. L. Bell and A. B. Slomson. *Models and Ultraproducts: An Introduction*. North-Holland, 1971.

[21] N. D. Belnap. A Useful Four-Valued Logic. In J. M. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37. D. Reidel Pub., 1975.

[22] N. D. Belnap. How Computer Should Think. In G. Ryle, editor, *Contemporary Aspects of Philosophy*, pages 30–55. Oriel Press, 1976.

[23] N. D. Belnap. Rescher's Hypothetical Reasoning. In E. Sosa, editor, *The Philosophy of Nicholas Rescher: Discussion and Replies*, pages 19–28. D. Reidel Pub., 1979.

[24] N. D. Belnap. Conjunctive Containment. In J. Norman and R. Sylvan, editors, *Directions in Relevant Logic*, pages 145–156. Kluwer Academic Pub., 1989.

[25] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency Management and Prioritized Syntax-Based Entailment. In *Proceedings of the Thirteenth International Joint Conferences on Artificial Intelligence*, volume 1, pages 640–645, 1993.

[26] S. Benferhat, D. Dubois, and H. Prade. A Local Approach to Reasoning Under Inconsistency in Stratified Knowledge Bases. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty 95*, Lecture Notes in Artificial Intelligence 946, pages 36–43. Springer Verlag, 1995.

[27] S. Benferhat, D. Dubois, and H. Prade. How to Infer from Inconsistent Beliefs without Revising? In *Proceedings of the Foureenth International Joint Conferences on Artificial Intelligence*, pages 1449–1455, 1995.

[28] S. Benferhat, D. Dubois, and H. Prade. Reasoning in Inconsistent Stratified Knowledge Bases. In *International Symposium on Multiple-Valued Logic*, pages 189–184, 1996.

[29] S. Benferhat, D. Dubois, and H. Prade. Some Syntactic Approaches to the Handling of Inconsistent Knowledge Bases: a Comparative Study, Part I: the Flat Case. *Studia Logica*, 58 (1):17–45, 1997.

[30] S. Benferhat, D. Dubois, and H. Prade. An Overview of Inconsistency-tolerant Inferences In Prioritized Knowledge Bases. In D. Dubois, E. P. Klement, and H. Prade, editors, *Fuzzy Sets, Logics and Reasoning About Knowledge*, Applied Logic Series vol 15, pages 395–418. Kluwer Academic Pub., 1999.

[31] S. Benferhat, D. Dubois, and H. Prade. Some Syntactic Approaches to the Handling of Inconsistent Knowledge Bases: a Comparative Study Part 2: the Prioritized Case. In E. Orłowska, editor, *Logic at Work: Essays Dedicated to the Memory of Helen Rasiowa*, Studies in Fuzziness and Soft Computing vol 24, pages 437–511. Physica-Verlag, 1999.

[32] S. Benferhat and L. Garcia. A Local Handling of Inconsistent Knowledge and Default Bases. In A. Hunter and S. Parson, editors, *Applications of Uncertainty Formalisms*, Lecture Notes in Artificial Intelligence 1455, pages 325–353. Springer Verlag, 1998.

[33] S. Benferhat, J. Lang, D. Dubois, H. Prade, A. Saffiotti, and P. Smets. Reasoning Under Inconsistency Based on Implicitly-specified Partial Qualitative Probability Relations: a Unified Framework. In *Proceedings: 15th National Conference on Artificial Intelligence (AAAI-98) and 10th Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*. AAAI-MIT Press, 1998.

[34] P. Besnard and A. Hunter. Quasi-Classical Logic: Non-Trivializable Classical Reasoning From Inconsistent Information. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty 95*, Lecture Notes in Artificial Intelligence 946, pages 44–51. Springer Verlag, 1995.

[35] P. Besnard and A. Hunter. Introduction to Actual and Potential Contradictions. In P. Besnard and A. Hunter, editors, *Handbook of Defeasible Reasoning and Uncertain*

Information Volume 2, Reasoning wiht Actual and Potential Contradictions, pages 1–12. Kluwer Academic Pub., 1998.

[36] P. Besnard and A. Hunter. A Logic-based Theory of Deductive Arguments. *Artificial Intelligence*, 128:203–235, 2001.

[37] P. Besnard and T. H. Schaub. Circumscribing Inconsistency. In *Proceedings of the Fifteenth International Joint Conferences on Artificial Intelligence*, volume 1, pages 150–155, 1997.

[38] P. Besnard and T. H. Schaub. Signed Systems for Paraconsistent Reasoning. *Journal of Automated Reasoning*, 20:191–213, 1998.

[39] W. Bibel. Constraint Satisfaction from a Deductive Viewpoint. *Artificial Intelligence*, 35:401–413, 1988.

[40] E. Birnbaum and E. L. Lozinskii. The Good Old Davis-Putnum Procedure Helps Counting Models. *Journal of Artificial Intelligence Research*, 10:457–477, 1999.

[41] H. A. Blair and V. S. Subrahmanian. Paraconsistent Logic Programming. *Theoretical Computer Science*, 68 (2):135, 1989.

[42] G. Boole. *An Investigation of the Laws of Thought*. Dover Pub., 1854.

[43] G. Brewka, J. Dix, and K. Konolige. *Nonmonotonic Reasoning: An Overview*. CSLI, 1997.

[44] B. Brown and P. K. Schotch. Logic and Aggregation. *Journal of Philosophical Logic*, 28:265–287, 1999.

[45] J. P. Burgess. Relevance: a Fallacy? *Notre Dame Journal of Formal Logic*, 22 (2):97–104, 1981.

[46] J. P. Burgess. Common Sense and 'Relevance'. *Notre Dame Journal of Formal Logic*, 24 (1):41–53, 1983.

[47] J. P. Burgess. Read on Relevance: a Rejoinder. *Notre Dame Journal of Formal Logic*, 25 (3):217–223, 1984.

[48] B. F. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1990.

[49] J. Copeland. The Genesis of Possible Worlds Semantics. *Journal of Philosophical Logic*, 31:99–137, 2002.

[50] R. H. Cowen. Two Hypergraph Theorems Equivalent to BPI. *Notre Dame Journal of Formal Logic*, 31:232–215, 1990.

[51] R. H. Cowen. Hypergraph Satisfiability. *Reports on Mathematical Logic*, 25:113–117, 1991.

[52] R. H. Cowen. Some Connections Between Set Theory and Computer Science. In *Computational Logic and Proof Theory: Third Kurt Gödel Colloquium*, Lecture Notes in Computer Science 713, pages 14–22. Springer Verlag, 1993.

[53] N. C. A. da Costa. On the Theory of Inconsistent Formal System. *Notre Dame Journal of Formal Logic*, 15 (4):497–510, 1974.

[54] J. de Kleer. An Assumption-Based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[55] J. de Kleer. Extending the ATMS. *Artificial Intelligence*, 28:163–196, 1986.

[56] J. de Kleer. Focusing on Probable Diagnoses. In *Proceedings of the National Conference on Artificial Intelligence (AAAI 1991)*, pages 842–848. Morgan Kaufmann, 1991.

[57] J. de Kleer, A. K. Mackworth, and R. Reiter. Characterizing Diagnoses. In *Proceedings of the National Conference on Artificial Intelligence (AAAI 1990)*, pages 324–330. Morgan Kaufmann, 1990.

[58] J. de Kleer, A. K. Mackworth, and R. Reiter. Characterizing Diagnoses and Systems. *Artificial Intelligence*, 56 (2-3):197–222, 1992.

[59] H. Deutsch. *A Family of Conforming Relevant Logics*. PhD thesis, Department of Philosophy, University of California, Los Angeles, 1981.

[60] H. Deutsch. Paraconsistent Analytic Implication. *Journal of Philosophical Logic*, 13:1–11, 1984.

[61] H. Deutsch. Relevance and Conformity. *Notre Dame Journal of Formal Logic*, 26 (4):455–462, 1985.

[62] H. Deutsch. Relevance and First-Degree Entailments. *Logique Et Analyse*, 28:3–20, 1985.

[63] J. Doyle. A Truth Maintenance System. *Artificial Intelligence*, 12:231–272, 1979.

[64] J. M. Dunn. A Modification of Parry's Analytic Implication. *Notre Dame Journal of Formal Logic*, 13 (2):195–205, 1972.

[65] J. M. Dunn. Intuitive Semantics for First-Degree Entailments and 'Coupled Tree'. *Philosophical Studies*, 29:149–168, 1976.

[66] J. M. Dunn. Relevance Logic and Entailment. In D. M. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic Volume 3: Alternatives To Classical Logic*, pages 117–224. D. Reidel Pub., 1986.

[67] R. Fagin and J. Y. Halpern. Belief, Awareness and Limited Reasoning. *Artificial Intelligence*, 34:39–76, 1988.

[68] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A Nonstandard Approach to the Logical Omniscience Problem. *Artificial Intelligence*, 79:203–240, 1995.

[69] K. Fine. Analytic Implication. *Notre Dame Journal of Formal Logic*, 27 (2):169–179, 1986.

[70] M. C. Fitting, V. W. Marek, and M. Truszczyński. The Pure Logic of Necessitation. *Journal of Logic and Computation*, 2 (3):349–373, 1992.

[71] D. M. Gabbay. On Second Order Intuitionistic Propositional Calculus with Full Comprehension. *Archiv für Mathematische Logik und Grundlagenforschung*, 16:177–186, 1974.

[72] D. M. Gabbay. *Investigations in Modal and Tense Logics with Applications To Problems in Philosophy and Linguistic*. D. Reidel Pub., 1976.

[73] D. M. Gabbay. *Labelled Deductive Systems: Volume 1*, volume Oxford Logic Guides 33. Oxford University Press, 1996.

[74] D. M. Gabbay and A. Hunter. Making Inconsistency Respectable: a Logical Framework For Inconsistency in Reasoning, Part I - a Position Paper. In *Fundamentals of Artificial Intelligence Research*, Lecture Notes in Artificial Intelligence 535, pages 19–32. Springer Verlag, 1991.

[75] L.T.F. Gamut. *Logic, Language, and Meaning: Volume 1 and 2*. University of Chicago Press, 1991.

[76] M. R. Garey and D. S. Johnson. *Computers and Intractability: a Guide to the Theory of NP- Completeness*. W. H. Freeman and Co., 1979.

[77] J. Garson. Modularity and Relevant Logic. *Notre Dame Journal of Formal Logic*, 30 (2):207–223, 1989.

[78] I. P. Gent, P. Prosser, and T. Walsh. The Constrainedness of Search. under review for Journal of the ACM, 1999.

[79] R. Goldblatt. *Mathematics of Modality*. CSLI, 1993.

[80] J. Grant. Classifications for Inconsistent Theories. *Notre Dame Journal of Formal Logic*, 19 (3):435–444, 1978.

[81] T. Hailperin. Best Possible Inequalities for the Probability of a Logical Function of Events. *American Mathematical Monthly*, 72 (4):343–359, 1965.

[82] T. Hailperin. *Boole's Logic and Probability*. North-Holland, 1976.

[83] W. H. Hanson. First-Degree Entailments and Information. *Notre Dame Journal of Formal Logic*, 21 (4):695–671, 1980.

[84] J. Hintikka. Information, Deduction, and the A Priori. *Noûs*, 4 (2):135–152, 1970.

[85] J. Hintikka. On Semantic Information. In J. Hintikka and P. Suppes, editors, *Information and Inference*, pages 3–27. D. Reidel Pub., 1970.

[86] J. Hintikka. Surface Information and Depth Information. In J. Hintikka and P. Suppes, editors, *Information and Inference*, pages 263–297. D. Reidel Pub., 1970.

[87] J. F. Horty. Nonmonotonic Foundations for Deontic Logic. In D. Nute, editor, *Defeasible Deontic Logic*, Studies in Epistemology, Logic, Methodology, and Philosophy of Science vol 263, pages 17–44. Kluwer Academic Pub., 1997.

[88] G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. Methuen & Co., 1968.

[89] G. E. Hughes and M. J. Cresswell. *A Companion to Modal Logic*. Methuen & Co., 1984.

[90] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, 1996.

[91] A. Hunter. Reasoning with Contradictory Information Using Quasi-Classical Logic. *Journal of Logic and Computation*, 10 (5):677–703, 2000.

[92] A. Hunter. A Semantic Tableau Version of First-Order Quasi-Classical Logic. In S. Benferhat and P. Besnard, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 6th European Conference, ECSQARU 2001, Toulouse, France, September 19-21, 2001, Proceedings*, Lecture Notes in Artificial Intelligence 2143, pages 544–555. Springer Verlag, 2001.

[93] A. Hunter and B. Nuseibeh. Analysing Inconsistent Specifications. In *Proceedings of the Third IEEE International Symposium on Requirements Engineering (RE'97)*, pages 78–86. IEEE Computer Society Press, 1997.

[94] A. Hunter and B. Nuseibeh. Managing Inconsistent Specifications: Reasoning, Analysis and Action. *ACM Transactions on Software Engineering and Methodology*, 7 (4):335–367, 1998.

[95] A. Hunter and S. Parson, editors. *Applications of Uncertainty Formalisms*. Lecture Notes in Artificial Intelligence 1455. Springer Verlag, 1998.

[96] S. Jaśkowski. On the Rules of Suppositions in Formal Logic. *Studia Logica*, 1, 1934. Reprinted in S. McCall (1967) *Polish Logic 1920–1939*, Oxford University Press, pp. 232–258.

[97] S. Jaśkowski. Propositional Calculus for Contradictory Deductive Systems. *Studia Logica*, 24:143–157, 1969.

[98] J. O. M. Jaspars. Logical Omniscience and Inconsistent Belief. In M. de Rijke, editor, *Diamonds and Defaults*, pages 129–146. Kluwer Academic Pub., 1993.

[99] R. E. Jennings, C. W. Chan, and M. J. Dowad. Generalised Inference and Inference Modelling. In *Proceedings of the Twelfth International Joint Conferences on Artificial Intelligence*, volume 2, pages 1046–1051, 1991.

[100] R. E. Jennings and P. K. Schotch. Some Remarks on (Weakly) Weak Modal Logics. *Notre Dame Journal of Formal Logic*, 22:309–314, 1981.

[101] R. E. Jennings and P. K. Schotch. The Preservation of Coherence. *Studia Logica*, 43 (1-2):89–106, 1984.

[102] D. K. Johnston. A Generalized Relational Semantics for Modal Logic. Master's thesis, Department of Philosophy, Simon Fraser University, 1978.

[103] B. Jónsson and A. Tarski. Boolean Algebras with Operators: Part I. *American Journal of Mathematics*, 73:891–939, 1951.

[104] B. Jónsson and A. Tarski. Boolean Algebras with Operators: Part II. *American Journal of Mathematics*, 74:127–162, 1952.

[105] G. Jumarie. *Relative Information: Theories and Applications*. Springer Verlag, 1990.

[106] H. Kautz and B. Selman. A General Framework for Knowledge Compilation. In H. Richter and M. Richter, editors, *Proceedings of International Workshop on Processing Delarative Knowledge*, Lecture Notes in Artificial Intelligence 567, pages 287–300. Springer Verlag, 1991.

[107] H. Kautz and B. Selman. Knowledge Compilation and Theory Approximation. *Journal of the ACM*, 43 (2):193–224, 1996.

[108] M. Kifer and E. L. Lozinskii. A Logic for Reasoning with Inconsistency. *Journal of Automated Reasoning*, 9:179–215, 1992.

[109] K. M. Knight. Measuring Inconsistency. *Journal of Philosophical Logic*, 31:77–98, 2002.

[110] K. M. Knight. Two Information Measures for Inconsistent Sets. *Journal of Logic, Language, and Information*, 12:227–248, 2003.

[111] K. M. Knight. Two Probabilistic Entailment and a Non-Probabilistic Logic. *Journal of the Interest Group in Pure and Applied Logics*, 11 (3):353–365, 2003.

[112] A. Kolany. Satisfiability on Hypergraphs. *Studia Logica*, 52 (3):393–404, 1993.

[113] A. Kolany. Consequence Operations Based on Hypergraph Satisfiability. *Studia Logica*, 58 (2):261–272, 1997.

[114] B. Konikowska. Rasiowa-Sikorski Deduction System: a Handy Tool for Computer Science Logics. In *Proceedings WADT's 98*, Lecture Notes in Computer Science 1589, pages 183–197. Springer Verlag, 1999.

[115] B. Konikowska and M. Białasik. Reasoning with First Order Nondeterministic Specifications. *Acta Informatica*, 36:375–403, 1999.

[116] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence*, 44:167–207, 1990.

[117] P. Lavers. Relevance and Disjunctive Syllogism. *Notre Dame Journal of Formal Logic*, 29 (1):34–44, 1988.

[118] E. J. Lemmon. *Beginning Logic*. Hackett Pub., 1978.

[119] H. J. Levesque. Logic and the Complexity of Reasoning. *Journal of Philosophical Logic*, 17 (4):355–389, 1988.

[120] C. I. Lewis and C. H. Langford. *Symbolic Logic*. New York, 1932.

[121] E. L. Lozinskii. Counting Propositional Models. *Information Processing Letters*, 41:327–332, 1992.

[122] E. L. Lozinskii. Information and Evidence in Logic Systems. *Journal of Experiment and Theoretical Artificial Intelligence*, 6:163–193, 1994.

[123] E. L. Lozinskii. Resolving Contradictions: a Plausible Semantics For Inconsistent Systems. *Journal of Automated Reasoning*, 12:1–31, 1994.

[124] J. J. Lu and E. Rosenthal. Annotations, Signs, and Generally Paraconsistent Logics. In E. A. Yfantis, editor, *Intelligent Systems: Third Golden West International Conference, Edited and Selected Papers Volume 1 and 2*, pages 143–157. Kluwer Academic Pub., 1995.

[125] A. K. Mackworth. The Logic of Constraint Satisfaction. *Artificial Intelligence*, 58:3–20, 1992.

[126] V. W. Marek, J. Treur, and M. Truszczyński. Representation Theory for Default Logic. *Annals of Mathematics and Artificial Intelligence*, 21 (2-4):343–358, 1997.

[127] F. Massacci. K-Clusters Tableaux: A Tool for Modal Logics and Inconsistent Belief Sets. *AI*IA Notizie*, 4:23–33, 1994.

[128] J. McCarthy. Circumscription – A Form of Nonmonotonic Reasoning. *Artificial Intelligence*, 13 (1-2):27–39, 1980.

[129] R. K. Meyer and R. Routley. Algebraic Analysis of Entailment. *Logique Et Analyse*, 15:407–428, 1972.

[130] C. Mortensen. Reply to Burgess and Read. *Notre Dame Journal of Formal Logic*, 27 (2):195–200, 1986.

[131] B. Nebel. Belief Revision and Default Reasoning: Syntax-Based Approaches. In J. A. Allen, R. Fikes, and E. Sandewall, editors, *Principle of Knowledge Representation and Reasoning: Proceedings of the Second International Conference (KR91)*, pages 417–428. Morgan Kaufmann, 1991.

[132] B. Nebel. Syntax-Based Approaches to Belief Revision. In P. Gärdenfors, editor, *Belief Revision*, pages 52–88. Cambridge University Press, 1992.

[133] A. Newell and H. Simon. GPS, a Program That Simulates Human Thought. In E. Feigenbaum and J. Feldman, editors, *Computers and Thought*, pages 279–296. McGraw-Hill, 1963.

[134] N. J. Nilsson. Probabilistic Logic. *Artificial Intelligence*, 28:71–87, 1986.

[135] W. T. Parry. Ein Axiomsystem für eube beye Art von Implikation (analytische Implikation). *Ergebrisse eines Mathematischen Colloquiums*, 4:5–6, 1933.

[136] S. Parsons and S. Green. Argumentation and Qualitative Decision Making. In A. Hunter and S. Parsons, editors, *Symbolic and Quantitative Approaches to Reasoning and Uncertainty, European Conference, ECSQARU'99, London, UK, July 1999. Proceedings*, Lecture Notes in Artificial Intelligence 1638, pages 328–339. Springer Verlag, 1999.

[137] F. J. Pelletier. A Brief History of Natural Deduction. *History and Philosophy of Logic*, 20:1–31, 1999.

[138] G. Priest. Minimally Inconsistent LP. *Studia Logica*, 50 (2):321, 1991.

[139] G. Priest. Negation as Cancellation and Connexive Logic. *Topoi*, 18:141–148, 1999.

[140] G. Priest. Paraconsistent Logic. In D. M. Gabbay, editor, *Handbook of Philosophical Logic Volume 6*. Kluwer Academic Pub., second edition, 2002.

[141] W. V. O. Quine. *Methods of Logic*. Henry Holt & Co., 1950.

[142] W. V. O. Quine. The Problem of Simplifying Truth Functions. *American Mathematical Monthly*, 59:521–531, 1952.

[143] W. V. O. Quine. A Way To Simplify Truth Functions. *American Mathematical Monthly*, 62:627–631, 1955.

[144] W. V. O. Quine. On Cores and Prime Implicants of Truth Functions. *American Mathematical Monthly*, 66:755–760, 1959.

[145] A. Ramesh, G. Becker, and N. V. Murray. CNF and DNF Considered Harmful for Computing Prime Implicants/Implicates. *Journal of Automated Reasoning*, 18:337–356, 1997.

[146] A. Ramesh, B. Beckert, R. Hähnle, and N. V. Murray. Fast Subsumption Checks Using Anti-Links. *Journal of Automated Reasoning*, 18:47–83, 1997.

[147] A. G. Ramesh. *Some Applications of Non Clausal Deduction*. PhD thesis, Department of Computer Science, State University of New York at Albany, 1995.

[148] V. Rantala. Impossible Worlds Semantics and Logical Omniscience. *Acta Philosophica Fennica*, 35:107–115, 1982.

[149] H. Rasiowa and R. Sikorski. *The Mathematics of Metamathematics*. PWN Warsaw, 1963.

[150] S. Read. Burgess on Relevance: a Fallacy Indeed. *Notre Dame Journal of Formal Logic*, 24 (4):473–481, 1983.

[151] R. Reiter. A Logic for Default Reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[152] N. Rescher. *Hypothetical Reasoning*. North-Holland, 1964.

[153] N. Rescher. *The Coherence Theory of Truth*. Oxford University Press, 1973.

[154] N. Rescher. *Plausible Reasoning*. Van Gorcum & Comp., 1976.

[155] N. Rescher and R. Brandom. *The Logic of Inconsistency: a Study in Non-Standard Possible World Semantics and Ontology*. American Philosophical Quarterly, 1979.

[156] N. Rescher and R. Manor. On Inference from Inconsistent Premisses. *Theory and Decision*, 1:179–217, 1970.

[157] R. Routley and R. K. Meyer. The Semantics of Entailment. In H. Leblanc, editor, *Truth, Syntax and Modality: Proceedings of the Temple University Conference on Alternative Semantics*, pages 199–243. North-Holland, 1973.

[158] R. Routley, V. Plumwood, R. K. Meyer, and R. T. Brady. *Relevant Logics and Their Rivals Volume 1*. Ridgeview Publishing Co., 1982.

[159] M. H. Saeedi and J. A. A. Sillince. Incorporating Rhetorical and Plausible Reasoning in a System for Simulating Argument. *Knowledge-Based Systems*, 12:113–127, 1999.

[160] H. Sawamura and S. Maeda. An Argumentation-Based Model of Multi-Agent Systems. preprint, Department of Information Engineering and Graduate School of Science and Technology, Niigata University, Japan, 2000.

[161] H. Sawamura, Y. Umeda, and R. K. Meyer. Computational Dialectics for Argument-based Systems. accepted for presentation at the Fourth International Conference on MultiAgents Systems (ICMAS 2000), Bosten, USA, 2000.

[162] M. Schaerf and M. Cadoli. Tractable Reasoning via Approximation. *Artificial Intelligence*, 74:249–310, 1995.

[163] T. H. Schaub. The Family of Default Logics. In P. Besnard and A. Hunter, editors, *Handbook of Defeasible Reasoning and Uncertain Information Volume 2, Reasoning wiht Actual and Potential Contradictions*, pages 77–134. Kluwer Academic Pub., 1998.

[164] P. K. Schotch and R. E. Jennings. Inference and Necessity. *Journal of Philosophical Logic*, 9:327–340, 1980.

[165] P. K. Schotch and R. E. Jennings. Modal Logic and the Theory of Modal Aggregation. *Philosophia*, 9:265–278, 1980.

[166] P. K. Schotch and R. E. Jennings. Non-Kripkean Deontic Logic. In R. Hilpinen, editor, *New Studies in Deontic Logic: Norms, Actions, and the Foundations of Ethics*, pages 149–162. Reidel Pub Co., 1981.

[167] P. K. Schotch and R. E. Jennings. On Detonating. In G. Priest, R. Routley, and J. Norman, editors, *Paraconsistent Logic: Essays on the Inconsistent*, pages 306–327. Philosophia Verlag, 1989.

[168] D. Scott. Completeness and Axiomatizability in Many-Valued Logic. In L. Henkin et al., editors, *Proceedings of Symposia in Pure Mathematics Volume XXV: Proceedings of the Tarski Symposium*, pages 411–435, 1974.

[169] C. E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27:379–423 and 623–656, 1948.

[170] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.

[171] J. K. Slaney. A General Logic. *Australasian Journal of Philosophy*, 68 (1):74–88, 1990.

[172] J. K. Slaney. **SCOTT**: A Model-Guided Theorem Prover. In *Proceedings of the Thirteenth International Joint Conferences on Artificial Intelligence*, volume 1, pages 109–115, 1993.

[173] J. K. Slaney, E. Lusk, and W. W. McCune. **SCOTT**: Semantically constrained **OTTER**. In A. Bundy, editor, *Proceedings of the 12th International Conference on Automated Deduction*, Lecture Notes in Artificial Intelligence 814, pages 764–768. Springer Verlag, 1994.

[174] R. M. Smullyan. *This Book Needs No Title: A Budget of Living Paradoxes*. Simon & Schuster, Inc., 1980.

[175] P. Suppes. *Introduction to Logic*. Van Nostrand/Reinhold Press, 1957.

[176] R. Sylvan. Relevant Containment Logics and Certain Frame Problems of AI. *Logique Et Analyse*, 31:11–24, 1988.

[177] Y. Umeda and H. Sawamura. Towards an Argument-based Agent System. In *Proceedings of Third International Conference on Knowledge-Based Intelligent Intelligent Information Engineering Systems*, pages 30–33. IEEE Press, 1999.

[178] A. Urquhart. A Semantic Theory of Analytical Implication. *Journal of Philosophical Logic*, 2:212–219, 1973.

[179] G. Wagner. *Vivid Logic: Knowledge-Based Reasoning with Two Kinds of Negation*. Lecture Notes in Artificial Intelligence 764. Springer Verlag, 1994.

[180] G. R. Wheeler. Kinds of Inconsistency. In *Paraconsistency: The Logical Way to the Inconsistent, Proceedings of the Second World Congress on Paraconsistency (WCP 2000)*, pages 511–522. Marcel Dekker, Inc., 2000.

[181] P. Wojtylak. 2 Sat is not equivalent to Boolean Prime Ideal Theorem. In E. Orłowska, editor, *Logic at Work: Essays Dedicated to the Memory of Helen Rasiowa*, Studies in Fuzziness and Soft Computing vol 24, pages 580–583. Physica-Verlag, 1999.

[182] P. Wong. Paraconsistent Inference and Preservation. Workshop on Logic in Computing Science, University of Technology, Sydney, 1998.