

AN ESSAY ON
CONTRACTION

—

André Fuhrmann

Studies in Logic, Language, and Information

Studies in Logic, Language and Information

Studies in Logic, Language and Information is the official book series of the European Association for Logic, Language and Information (FoLLI).

The scope of the series is the logical and computational foundations of natural, formal, and programming languages, as well as the different forms of human and mechanized inference and information processing. It covers the logical, linguistic, psychological and information-theoretic parts of the cognitive sciences as well as mathematical tools for them. The emphasis is on the theoretical and interdisciplinary aspects of these areas.

The series aims at the rapid dissemination of research monographs, lecture notes and edited volumes at an affordable price.

Managing editor: Robin Cooper, University of Gothenburg

Executive editor: Maarten de Rijke, University of Warwick

Editorial board:

Peter Aczel, Manchester University

Nicholas Asher, The University of Austin, Texas

Jon Barwise, Indiana University, Bloomington

John Etchemendy, CSLI, Stanford University

Dov Gabbay, Imperial College, London

Hans Kamp, Universität Stuttgart

Godehard Link, Universität München

Fernando Pereira, AT&T Bell Laboratories, Murray Hill

Dag Westerståhl, Stockholm University

An Essay on Contraction

André Fuhrmann

CSLI Publications

Center for the Study of Language and Information
Stanford, California

&

FoLLI

The European Association for
Logic, Language and Information

Copyright ©1997
CSLI Publications
Center for the Study of Language and Information
Leland Stanford Junior University
Printed in the United States
01 00 99 98 97 5 4 3 2 1

Library of Congress Cataloging-in-Publication Data

Fuhrmann, A. (André). 1958—
An essay on contraction / André Fuhrmann.
p. cm.

Includes bibliographical references (p. 107).

ISBN 1-57586-055-4 (hardcover : alk. paper).

ISBN 1-57586-054-6 (pbk. : alk. paper)

1. Belief and doubt. 2. Logic. I. Title.

BC199.B4B84 1996

121.6-dc20

96-9068

CIP

Ⓢ The acid-free paper used in this book meets the minimum requirements of the American National Standard for Information Sciences—Permanence of Paper for Printed Library Materials, ANSI Z39.48-1984.

CSLI Publications reports new developments in the study of language, information, and computation. In addition to lecture notes, our publications include monographs, working papers, revised dissertations, and conference proceedings. Our aim is to make new results, ideas, and approaches available as quickly as possible. Please visit our website at <http://csli-www.stanford.edu/publications/> for comments on this and other titles, as well as for changes and corrections by the author and publisher.

Contents [\(clickable\)](#)

| | |
|---|-----|
| <i>Preface</i> | v |
| 1. <i>Introduction</i> | 1 |
| 2. <i>Theories and Theory Change</i> | 8 |
| 2.1 Sentential language and consequence operations | 9 |
| 2.2 Three kinds of theory change | 17 |
| 2.3 The AGM postulates | 20 |
| 2.4 Constructions of contractions | 27 |
| 3. <i>General Contraction</i> | 33 |
| 3.1 An imbalance | 34 |
| 3.2 Closure | 35 |
| 3.3 Very small changes | 37 |
| 3.4 Do, need and should we multiply contract? | 37 |
| 3.5 Package contraction and choice contraction | 42 |
| 3.6 Basic properties of multiple contractions | 45 |
| 3.7 Tools for constructing general contraction functions | 54 |
| 3.8 Constructing contractions | 55 |
| 3.9 Reduction strategies revisited | 61 |
| 3.10 Theory contraction through base replacement | 64 |
| 3.11 Why general contraction is general enough | 67 |
| 3.12 Logical subtraction revisited | 70 |
| 4. <i>Revision and Merge</i> | 79 |
| 4.1 Revision | 80 |
| 4.2 Merge | 87 |
| 4.3 Inference | 92 |
| 4.4 Expectation inference | 93 |
| 4.5 Merge inference | 97 |
| 5. <i>Everything in Flux: Dynamic Ontologies</i> | 105 |
| 5.1 Closure systems and ontological spaces | 105 |
| 5.2 Travelling in ontological space | 107 |
| 5.3 Finding (good) examples | 108 |
| <i>References</i> | 114 |
| <i>Notation</i> | 121 |

*For
Richard Sylvan*

Preface

This essay has grown out of a series of investigations aimed at generalising and extending the logic of belief (or theory) change in the sense introduced by Carlos Alchourrón, Peter Gärdenfors and David Makinson in a series of seminal papers. Some of the generalisations and extensions were already hinted at in the author's PhD thesis (1988). First results were published in (1991) and in a joint paper with Sven Ove Hansson (1994).

The research reported here was conducted while I held a research position at the University of Konstanz. I am grateful to Jürgen Mittelstraß for creating the academic environment in which my study of contractions flourished. Most of the essay was written while visiting Indiana University in Bloomington (August to October 1993) and Columbia University (February to April 1994) with a grant from the German Research Council (DFG). I thank my hosts, J. Michael Dunn and Isaac Levi, for their hospitality. In both universities I had many good discussions; apart from Mike Dunn and Isaac Levi I wish to mention and thank Anil Gupta and Raymundo Morrado in Bloomington, and Horacio Arló Costa and John Collins in New York.

Thanks are also due to the theory-changers community for many discussions over the years: apart from those just mentioned, Peter Gärdenfors, Sven Ove Hansson, David Makinson, Wlodek Rabinowicz, and others. I particularly thank Sven Ove Hansson for engaging in joint research with me and for kindly permitting that some of our jointly reached results be reproduced in Chapter Three.

Finally, I am grateful to two anonymous referees for the publisher who pointed out faults and gaps and helped improving the essay in many ways.

1

Introduction

William James describes the process by which we acquire new beliefs as follows:

“The process [by which any individual settles into new opinions] is always the same. The individual has a stock of old opinions already, but he meets a new experience that puts them to a strain. [...] The result is an inward trouble to which his mind till then had been a stranger, and from which he seeks to escape by modifying his previous mass of opinions. He saves as much of it as he can, for in this matter of belief we are all extreme conservatives. So he tries to change first this opinion, and then that (for they resist change very variously), until at last some new idea comes up which he can graft upon the ancient stock with a minimum of disturbance of the latter, some idea that mediates between the stock and the new experience and runs them into one most felicitously and expediently.

“The new idea is then adopted as the true one. It preserves the older stock of truths with a minimum of modification, stretching them just enough to make them admit the novelty, but conceiving that in ways as familiar as the case leaves possible. [...] New truth is always a go-between, a smoother-over of transitions. It marries old opinions to new facts so as ever to show a minimum of jolt, a maximum of continuity.” (1907, 59ff.)

The purpose of this essay is to probe deeper into the formal structure of the process described informally in the passage just cited. For formal

structure there is, and the structure is both simple and vexing at the same time.

There are beliefs and there is a relation that compares their resistance to change. New beliefs may either fit in smoothly with old beliefs or they may “put them to a strain”. In the latter case the old beliefs together with the new beliefs form an *incoherent* whole. So here are three elements of a formal analysis:

- (i) a domain of *beliefs*,
- (ii) a relation of *comparative resistance* between beliefs, and
- () a property of stocks of beliefs, called *incoherence*.

(The latter is not tagged with a numeral because we shall replace it in a moment with a better understood notion.)

Given the current stock of beliefs B and a potential new belief p such that B and p cannot be believed together on pain of incoherence, the task is to somehow use the resistance relation to determine the maximum sub-collection of B that we may allow to survive the pressure to accommodate p .

A common abstraction of “collections” of beliefs is to view them as *sets*. “Living” sets of belief have no doubt a much richer structure. For example, we acquire beliefs in a certain temporal order and from many different sources. Many beliefs are also tagged with an emotional colouring. As far as these aspects matter for how we should change our beliefs, they can be taken account of by the notion of comparative resistance to change. If “fresh” (or “long-held”, or “warm”) beliefs are to have priority over “stale” (or over “untried”, or over “cold”) ones, let them be more resistant; similarly with beliefs from more or less reliable sources. So sets-*cum*-resistance seems to provide a representation of belief states that is adequate for the present purpose.

Next we turn to coherence. This notion is intimately linked to that of consistency. However, not all forms of incoherence are outright cases of inconsistency. For example, the belief that a particular coin is fair is consistent with the belief that the next thousand tosses of that coin will land head up. But, so one wishes to say, the two beliefs are obviously at odds with each other.

The oddity can be explained by pointing out that inductive inference from the belief that the coin is fair would lead to a conclusion which is inconsistent with the further belief that the next thousand tosses will land head up. We must be prepared to draw such inductive inference, for otherwise—if, say, we know that the fabulous Ricky Jay will toss the coin—the incoherence evaporates. Thus the kind of incoherence exhibited in this case can be

reduced to the well-understood notion of (logical) consistency.

There is another sense of incoherence which seems more remote from an analysis in terms of consistency. Consider the belief that for every prime number $n < 1000$, the n th toss of a particular fair coin will land head up. Again, this is an odd belief and it may well be called incoherent with the belief that the coin is fair. But unlike in the former case, inductive inference will uncover no inconsistency; for such inference will yield no predictions for particular tosses. How then should the oddity be accounted for?

This is certainly an intriguing question and I recommend that the interested reader turns to Lycan (1988, ch. 8) for an answer. Fortunately, we do not here need to pursue the question. For incoherence in this last, weaker sense does not call for internal adjustments.

To see why, it is best to continue our example. Suppose that you believe that the coin is fair and that testimony of the senses on a thousand trials convinces you that the coin lands head up at every prime-numbered toss. Surely, this is odd. But is it incoherent in the sense that you cannot combine the two beliefs? What else is there to do but to simply take the new evidence on board and wait for future experiments to provide some explanation of what has been witnessed? Incoherence in the third, weak sense, though generating cognitive stress, never calls for immediate adjustments; instead it demands further investigation to relieve the stress.

We therefore suggest that the sort of incoherence that induces belief change can be reduced to inconsistency (before or after induction). In any case, our investigation will cover only those changes that are initiated by incoherencies that can be analysed in terms of inconsistency. The domain of change thus covered is large and important enough; the question as to just how large it is, need not delay us here.

Consistency, in turn, is a property interdefinable with that of logical consequence: two beliefs can be held consistently just in case the one does not follow from the negation of the other. Hence, the third ingredient in a general theory of belief change is

- (iii) a relation of logical *consequence*.

(It will emerge that we shall need elementary properties only of the consequence relation. In particular, we shall not assume that the consequence relation is codified by some particular logic, classical, intuitionist, or other.)

Now that we have introduced some of the abstractions underlying our investigation of belief change, let us turn to the kinds of belief changes that may occur.

Given a decision to accommodate new information, two possibilities can arise: either the new information is consistent with what is already believed

or it is not.

In the first case no re-arrangement of old beliefs is necessary. The new information can be adopted into our stock of beliefs without any disturbance to the latter. This kind of change is simply additive (or “monotonic”) and, as we shall see shortly, it can be easily characterised.

In the second case incisions into old beliefs are called for. Before the new piece of information can be accepted as belief, the old stock of belief has to be altered so as to make room for consistently receiving the new belief. In the present case then, belief change is a two-steps process: first the old stock of beliefs must be contracted so as to open one’s mind towards the new candidate belief; then the new belief may be added without threat of inconsistency.

It is not difficult to see that subtracting a belief from a body of opinions is not usually a straightforward matter. For a stock of beliefs is not simply an odd bag of isolated opinions. Beliefs are usually part of a more or less tight web; the knots are enmeshed with each other by various relations of support. (Later we shall model support by means of logical consequence.) With many beliefs it is the case that they are supported by a set of other beliefs without being supported by any particular member of that set. When such an essentially multiply supported belief is to be retracted, at least one of the supporting beliefs needs to be removed too; otherwise the belief to be retracted will be reinstated. This situation immediately raises the question which one of the supporting beliefs to remove. Contraction thus involves an element of choice.

Given that the mere addition of beliefs is a relatively unproblematic process, contractions will occupy centre stage in a theory of belief change and they will accordingly form the central theme of this essay.

The quotation from James—and our attempt at distilling from it the outlines of a formal theory—may create the impression that the aim is a descriptive rather than a prescriptive theory of belief change. This impression would be mistaken, however. James very clearly held the view that insofar as we modify our beliefs in the described manner, we act rationally. (Incidentally, he also believed that, by and large, we act as we should; hence, the prescriptive theory happens to be also descriptive—by and large.)

The basis for this view is the seminal *belief-doubt model* of inquiry first put forward by Peirce (1877). This model is a unifying thread in the doctrines of otherwise differing philosophers of pragmatic persuasions, like Peirce, James, Dewey, Schiller, Ramsey, or Quine. Among contemporary philosophers Isaac Levi (1967, 1980, 1991) has developed the model in great detail. Levi takes particularly seriously Peirce’s idea that inquiry must start

from a corpus of *full* beliefs.

According to the belief-doubt model each inquiry must proceed from a given corpus of beliefs which are taken to be true as long as no evidence to the contrary is encountered. The inquirer suspends or retracts a belief only if weighing the evidence forces her to conclude that holding on to the belief in question incurs too high a risk of error. Cognitive conservatism is an essential part of this model. My current beliefs have, presumably, guided me well in previous inquiries; they have so far stood the test of time; they have proved reliable. I can do no better than (fully) accept what has proved reliable and should only retract when forced to do so. When forced to retract, I should do so by performing only a minimal incision into what I must take to be a largely reliable body of doctrine.

This model stands in stark contrast to the Cartesian *doubt-belief model* according to which inquiry should start with summary doubt: suspend all belief until secure foundations have been laid. It would be tedious to rehearse here the modern critique of the Cartesian model. I must leave it to the reader to form her own opinion as to how much of the formal theory to be presented in the next chapters makes sense in the context of a Cartesian epistemology. On the one hand, the theory of contractions will allow for singling out a distinguished subset of a set of beliefs. This subset may well be assumed to represent the foundations from which other beliefs can be justified by inference. On the other hand such foundational beliefs may be retracted; hence, they cannot carry Cartesian certainty. Moreover, it seems difficult to accommodate within a Cartesian framework our concern with *minimising* incisions into beliefs in the face of recalcitrant evidence—should the Cartesian not follow a course of *maximal* retractions?

A belief that is adopted as the result of engaging in Jamesian (or Peircean) belief changes, is, according to advocates of the belief-doubt model, a belief as reliable as one can ever be obtained. In that sense the belief may be called *justified*. Thus, James' theory is not a piece of armchair psychology but a contribution to epistemology: that part of philosophy that is concerned with how beliefs and belief changes score honorific titles, like “justified” or “rational” or perhaps even “true” respectively “truth-conducive”. An exposition of the formal backbone of such a theory should accordingly also be viewed as a contribution to epistemology.

* * *

No value can be added to an intellectual enterprise simply by subsuming it under a recognised discipline. In that respect, the last remark of the preceding paragraph should be superfluous. Yet, for various reasons,

many philosophers shy away from formal methods. This reflex sometimes degenerates into a sentiment *against* formal methods, particularly in those areas of philosophy that are considered to be at the heart of the discipline, such as metaphysics, ethics or epistemology. Understandable as the reflex may be, the sentiment has no doubt a harmful effect on philosophy. For it discourages inquiry where it should be pursued and it erects barriers to neighbouring disciplines where there should be bridges.

One discipline, artificial intelligence (AI), deserves particular mention here. AI has always been impressed with the flexibility of the human mind: with the ability to absorb new information quickly and to jump on the whole successfully to conclusions from deductively insufficient premisses. In what follows, a contribution towards an account of these abilities will be made. We shall see what minimal structure we must assume if new information is to be absorbed while keeping low the risk of error. The theory of belief change yields also modellings of various kinds of non-deductive, defeasible inference, some of them will be detailed in Chapter Four.

The formal nature of these theories is essential for the bridge to AI. For at the heart of AI is a computer-model of the mind, and computers like to be talked to formally. Even if the talk of a “computer-model of the human mind” is construed very widely—e.g. admitting presently not realisable architectures—it remains true that formal theories of the kind presented in this essay are particularly helpful to AI: they are free from assumptions about the medium carrying the formal structure; their elements are well-understood from a computational point of view; and they admit of a number of equivalent presentations which may however differ in their computational complexity.

* * *

The present essay is largely self-contained. In Chapter Two we present the background to the subsequent studies. Formal languages and an abstract notion of logical consequence, due to Alfred Tarski, will be introduced. Then we proceed to outlining the classical logic of theory change, developed by Carlos Alchourrón, Peter Gärdenfors and David Makinson. The classical theory will subsequently be extended and revised and various ways.

Chapter Three contains the general theory of contractions. Different kinds of contraction operations are characterised both axiomatically and by means of explicit constructions. Representation results to the effect that axioms and constructions match will be proved.

Chapter Four focuses on two operations, revision and merge, which combine in different ways contractions and expansions. We link the two change

operations to two simple examples of nonmonotonic inference: revision to so-called expectation inference, and merge to a cancellation theory of negation leading to a paraconsistent inference relation.

The theory of contractions presented here is in no way essentially tied to contractions of systems of *belief*. All results carry over without much modification to arbitrary closure systems. Closure systems, being of a very general structure, are ubiquitous. Chapter Five opens the prospect for a much larger class of applications than originally envisaged by those who introduced and developed the theory of contractions with belief (or theory) change in mind.

2

Theories and theory change

In this chapter we introduce the formal background to most of our investigations. The concept of a consequence operation on the sentences of some regimented language was introduced by Alfred Tarski (1930). Underlying the logic of theory change is a very general concept of a theory (also introduced by Tarski) as a set of sentences containing all of its logical consequences. Indeed, the logic of theory change—as presented in this essay—is easily generalised to pertain to any closure system whatsoever. This generalisation is studied in Fuhrmann (1995, 199+a) under the heading of “dynamic ontologies”.

The results pertaining to consequence operations and their associated consequence relations are logical folklore; hence, no proofs will be stated here. For a comprehensive survey of results in this area the reader is referred to Wójcicki (1988).

After having provided the requisite background in logical matters, we then proceed to rehearsing the key definitions and results of the AGM (Alchourrón, Gärdenfors, Makinson) theory of theory change. Full proofs may be found in the writings of AGM; sometimes the required arguments will be sketched as a preparation for the results to be presented in subsequent chapters.

2.1 Sentential languages and consequence operations

Sentential languages. Our topic is theory change and we thus need some means of representing the objects of our investigation. We assume that theories are representable in some sentential language. A *sentential language*, as we shall use the term here, is built up from a set basic sentences (*atoms*) by applying certain sentence forming operations (*connectives*). For the most part of what follows, we do not need to make specific assumptions about the nature of the connectives involved. But, by default as it were, certain parts of the classical theory of theory change—the theory of Carlos Alchourrón, Peter Gärdenfors and David Makinson (AGM)—are built on the assumption that classical logic ought to govern the notion of a theory (a definition will be provided in a moment). Those readers who have no quarrel with that assumption may simply assume that whenever connectives are displayed, they have their Boolean meaning unless, of course, stated otherwise. Connectives that may thus be interpreted are:

| | |
|-------------------|-------------|
| \perp | falsum |
| \top | verum |
| \neg | negation |
| \wedge | conjunction |
| \vee | disjunction |
| \rightarrow | implication |
| \leftrightarrow | equivalence |

However, the connectives just mentioned do not need to be read with full Boolean force. Thus the more discerning—and, perhaps, less Boolean-minded—reader may want to take notice of the fact that only a few and mostly elementary properties of these connectives will ever be needed.

We use p, q, r, \dots to variably refer to the atoms of a sentential language (though there will be few occasions to explicitly mention atoms). Lower case Greek letters

$$\alpha, \beta, \gamma, \dots$$

will be used as variables ranging over sentences and capital Roman letters

$$A, B, C, \dots, X, Y, Z$$

will range over sets of sentences. Calligraphic letters,

$$\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots, \mathcal{X}, \mathcal{Y}, \mathcal{Z}$$

stand for collections of sets of sentences.

All kinds of variables may occur decorated with numerals or primes. Plain or decorated T is reserved as a variable for a special kind of set of sentences: theories. The set of all formulae of the language under consideration will be denoted by ' $Fm\mathcal{L}$ '.

Finally, it will be convenient to succinctly express that some set A is a *nonempty and finite subset* of some set B ; we shall write this as:

$$A \subseteq_f B$$

Consequences. A theory is not any odd set of sentences but one which bears certain marks of quality. The quality that will be of prime interest in the sequel is this: Any sentence that follows from a certain theory is part of that theory. This immediately raises the question, how to determine whether a sentence follows from a set of sentences. In general the answer seems to be clear. A sentence α follows from a set A just in case α can be derived from A by applying rules of inference taken from an antecedently agreed upon set of such rules. But not only do different logics differ on which rules of inference should be applied but the special sciences each supply their own special corpus of such rules.

To say something about theories in sufficient generality, we must thus take the relation of 'follows from' not as defined—using a specific set of rules—but as primitive and lay down a number of integrity constraints which any specific consequence relation should obey. This is the approach first followed by Alfred Tarski in his seminal paper 'Fundamental concepts of the methodology of the deductive sciences' (1930).

Tarski (1930) uses a consequence operation, Cn , as primitive rather than a consequence relation. The difference is mostly a matter of notation, as we shall see in a moment. Intuitively, $Cn(A)$ denotes the set of sentences that follow from the set A (the syntactic perspective) or that are true in all models of A (the semantic perspective). According to Tarski a *consequence operation* is any mapping of a set of sentences to a set of sentences satisfying the conditions of reflexivity, monotonicity and idempotence; in addition, finitary consequence operations are required to also satisfy a condition of finitude:

$$\begin{aligned} A &\subseteq Cn(A) && \text{(Refl)} \\ A \subseteq B &\implies Cn(A) \subseteq Cn(B) && \text{(Mon)} \\ Cn(Cn(A)) &\subseteq Cn(A) && \text{(Idm)} \\ Cn(A) &= \bigcup \{Cn(A') : A' \text{ is a finite subset of } A\} && \text{(Fin)} \end{aligned}$$

We shall refer to this set of postulates as the *Tarski axioms*.

Given a consequence operation, there are a number of ways to define a consequence relation. (The sense of multitude will be explained in the remarks following the observation below.) Two of these will be particularly useful in what follows.

DEFINITION 2.1 For sets of sentences A and B :

1. $A \vdash B \iff B \cap \text{Cn}(A) \neq \emptyset$;
2. $A \Vdash B \iff B \subseteq \text{Cn}(A)$.

The relation \vdash holds between A and B if some of B are among the consequences of A , and the relation \Vdash holds if all of B is contained in the consequences of A . Since we shall make frequent use of the two consequence relations \vdash and \Vdash , it will be worthwhile emphasizing that, according to the above definition, $A \not\vdash B$ is another way of saying that none of B is a consequence of A and $A \not\Vdash B$ means that B is not contained in the consequences of A .

We shall economise a little in writing down elements of the consequence relations: for instance $A \cup B \vdash C \cup \{\alpha\}$ will be abbreviated to $A, B \vdash C, \alpha$, and similarly for \Vdash .

As to the relation between \vdash and \Vdash note, first, that for singleton right-hand-sides the two relations coincide:

$$A \vdash \alpha \iff A \Vdash \alpha \quad (\iff \alpha \in \text{Cn}(A))$$

Note, second, that

$$A \Vdash B \ \& \ B \Vdash A \iff \text{Cn}(A) = \text{Cn}(B)$$

However, from $A \vdash B$ and $B \vdash A$ (notation: $A \dashv\vdash B$) it does not in general follow that $\text{Cn}(A) = \text{Cn}(B)$ —though the converse does of course hold as long as A and B are non-empty. For $\text{Cn}(A) = \text{Cn}(B)$ we shall also use the notation $A \equiv B$.

The following observation is due to Scott (1974, Theorem 1.2).

OBSERVATION 2.2 *If the operation Cn satisfies the Tarski axioms, then the relation \vdash of Definition 2.1.1 satisfies the Scott axioms:*

$$\begin{aligned} A \cap B \neq \emptyset &\implies A \vdash B && \text{(Overlap)} \\ A \subseteq A' \ \& \ B \subseteq B' \ \& \ A \vdash B &\implies A' \vdash B' && \text{(Weakening)} \\ A, \beta \vdash C \ \& \ A \vdash \beta, C &\implies A \vdash C && \text{(Cut)} \\ A \vdash B &\implies \exists \text{ finite } A' \subseteq A : A' \vdash B && \text{(Compactness)} \end{aligned}$$

PROOF. See Scott (1974, p.415). ■

Remarks. The consequence relation of Definition 2.1.1 is the minimal (compact) relation \vdash_{\min} of Scott (1974, Theorem 1.2). Many more consequence relations, i.e. relations satisfying the Scott axioms, can be derived from a given consequence operation. The family of such consequence relations is bounded by our relation and the maximal relation defined as

$$A \vdash_{\max} B \text{ iff } \forall A' \supseteq A: \bigcap_{\beta \in B} \text{Cn}(A' \cup \{\beta\}) \subseteq \text{Cn}(A')$$

The “pure” theory of consequence omits the finitude and compactness axioms. But they are usually an essential part of the “applied” theory and will be assumed to hold in what follows. (Scott omits compactness since it is trivially implied by his assumption—not made here—that \vdash obtains between *finite* sets only.)

We may also take the reverse route, i.e., taking the notion of a consequence relation (satisfying the Scott axioms) as primitive and then define an operation as follows:

$$\text{Cn}(A) = \{\beta: A \vdash \beta\}$$

The operation thus defined will be a consequence operation in the sense of Tarski.

The cut principle is stated in a weak, sentential version. This is all that is required for our purposes in later chapters. We note, however, that the Tarski axioms license a stronger version of cut, called ‘cut for sets’ in Shoesmith and Smiley (1974, Sec. 2.1):

$$\text{if } A, B_1 \vdash B_2, C \text{ for each partition } (B_1, B_2) \text{ of } B, \text{ then } A \vdash C$$

where B_1 and B_2 partition B just in case they are disjoint sets whose union is B . *(End of Remarks)*

Note that while sets occurring on the right-hand-side of the turnstile are interpreted “disjunctively”, they are so interpreted in a subtly different sense of the word than the one underlying, say, Gentzen’s notion of a sequent. The disjunctive interpretation underlying Gentzen’s sequents amounts to the following equivalence:

$$A \vdash \beta_1, \dots, \beta_n \iff A \vdash \beta_1 \vee \dots \vee \beta_n$$

In contrast, Definition 2.1 gives the equivalence

$$A \vdash \beta_1, \dots, \beta_n \iff A \vdash \beta_1 \text{ or } \dots \text{ or } A \vdash \beta_n$$

Under the weak assumption

$$A \vdash \beta_1 \text{ or } A \vdash \beta_2 \implies A \vdash \beta_1 \vee \beta_2 \quad (*)$$

about the behaviour of disjunction, the relation defined above is stronger than Gentzen's. The two relations only coincide under the excessively strong assumption that \vdash satisfies the strong disjunction property, expressed by the converse of (*).

Gentzen had, of course, very good reasons for proceeding as he did. For the present purpose, however, our consequence relations will work more smoothly—with one exception. In Chapter 4 we shall need to extend the operation of negating a sentence to cover also sets. An important constraint on any acceptable (for the purposes of Chapter 4) definition of set-negation is that the following equivalence holds:

$$A, B \vdash \perp \iff A \vdash \neg B \quad (\text{N})$$

(We use the symbol \neg to denote both ordinary sentential negation and the stretched notion of negation that also applies to sets.) Now, a pleasantly simple definition of set-negation would let the negation of a set be just the set of negations of all of its elements:

$$\neg A = \{\neg\alpha : \alpha \in A\} \quad *(\text{D}\neg)$$

This does work with Gentzen's consequence relation but not with ours. That is, (N) follows from $*(\text{D}\neg)$ according to Gentzen's understanding of \vdash but not according to ours. The problem is the left-to-right direction of (N). For a drastic counterexample let A be empty and let $B = \{p, \neg p\}$. Then, clearly, $p, \neg p \vdash \perp$ while neither $\emptyset \vdash p$ nor $\emptyset \vdash \neg p$.

The definition we require is slightly more complicated: the negation of a set A consists of the negated conjunctions of all finite nonempty subsets of A , i.e.

$$\neg A = \{\neg \bigwedge A' : A' \subseteq_f A\} \quad (\text{D}\neg)$$

Given minimal assumptions concerning conjunction and negation this does work as the next lemma records. The minimal assumptions needed are

$$\text{Cn}(\alpha, \beta) = \text{Cn}(\alpha \wedge \beta) \quad (\text{Conj})$$

$$\perp \in \text{Cn}(A \cup \{\alpha\}) \iff \neg\alpha \in \text{Cn}(A) \quad (\text{Neg})$$

LEMMA 2.3 *Let Cn satisfy the Tarski axioms as well as (Conj) and (Neg); let \vdash be defined as in Definition 2.1.1; and, for each set $B \subseteq \text{Fml}$, let $\neg B$ be defined according to (D \neg). Then the equivalence (N) holds.*

PROOF. (\implies) Suppose that $A, B \vdash \perp$. It follows by compactness (C) that there exists a finite subset B' of B such $A, B' \vdash \perp$. In view of (Conj)

we may represent B' by the conjunction $\bigwedge B'$. It follows by (Neg) from $A, \bigwedge B' \vdash \perp$ that $A \vdash \neg \bigwedge B'$. Since $B' \subseteq_f B$, we have $A \vdash \neg B$ as required.

(\Leftarrow) Note first that \vdash enjoys the property of *utmost compactness on the right*:

$$A \vdash B \implies \exists \beta \in B: A \vdash \beta \quad (*)$$

(This follows immediately from \vdash 's definition.) Suppose now that $A \vdash \neg B$. By (*) we have $A \vdash \neg \beta$, for some $\neg \beta \in \neg B$. So by (Neg), $A, \beta \vdash \perp$ whence, by weakening, $A, B, \beta \vdash \perp$. But since β is a conjunction of elements of B , $B \vdash \beta$; hence, by cut, $A, B \vdash \perp$. ■

In view of the following lemma a more economical definition of *finite* set negation is also possible.

LEMMA 2.4 *If B is a finite set, then $A \vdash \neg B$ if and only if $A \vdash \bigvee \{\neg \beta: \beta \in B\}$.*

No discussion of consequence would be complete without mentioning—along with Tarski and Scott—Bolzano. Bolzano (in his *Wissenschaftslehre* (1837), §155) defines a relation of logical consequence that extends to multiple conclusions. Interestingly, and in contrast to later developments, he opts for a “conjunctive” interpretation of the right-hand-side, thus, in effect, introducing our relation \Vdash of Definition 2.1.2. It is pleasing to observe that the agreement between Tarski and Scott extends to Bolzano:

OBSERVATION 2.5

1. *If the operation Cn satisfies the Tarski axioms, then the relation \Vdash of Definition 2.1.2 satisfies the Bolzano axioms:*

$$B \subseteq A \implies A \Vdash B \quad (\text{Inclusion})$$

$$A \Vdash B \implies A, C \Vdash B, C \quad (\text{Isotonicity})$$

$$A \Vdash B \ \& \ B \Vdash C \implies A \Vdash C \quad (\text{Transitivity})$$

$$\text{If } B \text{ is finite and } A \Vdash B \implies \exists \text{ finite } A' \subseteq A: A' \Vdash B \quad (\text{Compactness}^*)$$

2. *If the relation \Vdash satisfies the Bolzano axioms, then the operation Cn defined as $\text{Cn}(A) = \{\beta: A \Vdash \beta\}$, for each $A \in \text{Fml}$, satisfies the Tarski axioms.*

PROOF. *Ad 1:* (Inclusion) follows immediately from (Refl). For (Isotonicity) suppose $A \Vdash B$, i.e. $B \subseteq \text{Cn}(A)$. Then $B \cup C \subseteq \text{Cn}(A) \cup C$. Further, $\text{Cn}(A) \cup C \subseteq \text{Cn}(\text{Cn}(A) \cup C) \subseteq \text{Cn}(A \cup C)$ by (Refl) and (Idm). So $B \cup C \subseteq \text{Cn}(A \cup C)$. i.e. $A, C \Vdash B, C$. For (Transitivity) suppose $A \Vdash B$ and $B \Vdash C$, i.e. $B \subseteq \text{Cn}(A)$ and $C \subseteq \text{Cn}(B)$. Then $\text{Cn}(B) \subseteq \text{Cn}(\text{Cn}(A))$

by (Mono). So $C \subseteq \text{Cn}(\text{Cn}(A)) \subseteq \text{Cn}(A)$ by (Idm). For (Compactness*) assume that B is a finite subset of $\text{Cn}(A)$. Then it follows by (Fin) that

$$B \subseteq \bigcup \{ \text{Cn}(A') : A' \text{ is a finite subset of } A \}$$

Call the right-hand-side of this inclusion ‘ C ’. Since B is finite, $B = \{\beta_1, \dots, \beta_n\}$ whence $\beta_i \in C$ ($1 \leq i \leq n$), i.e. $\beta_i \in \text{Cn}(A'_i)$ for some finite subset A'_i of A . It follows that

$$\{\beta_1, \dots, \beta_n\} \subseteq \text{Cn}(A'_1) \cup \dots \cup \text{Cn}(A'_n)$$

with each A_i a finite subset of A . But

$$\text{Cn}(A'_1) \cup \dots \cup \text{Cn}(A'_n) \subseteq \text{Cn}(A_1 \cup \dots \cup A_n)$$

where $A' = A_1 \cup \dots \cup A_n$ is again finite. Hence, $B \subseteq \text{Cn}(A') \subseteq A$ with A' finite as required.

Ad 2: (Refl) is immediate by (Inclusion). For (Mon) assume $A \subseteq B$ and $A \Vdash \beta$. (It will suffice to show that $B \Vdash \beta$). Then $A, B \Vdash \beta, B$ by (Isotonicity) and since $A \subseteq B$, we have $B \Vdash \beta, B$. But $\beta, B \Vdash \beta$ by (Inclusion). Hence, $B \Vdash \beta$ by (Transitivity). To show (Idm) assume $\beta \in \text{Cn}(\text{Cn}(A))$. So $\{\alpha : A \Vdash \alpha\} \Vdash \beta$. But $A \Vdash \{\alpha : A \Vdash \alpha\}$ whence $A \Vdash \beta$ by (Transitivity). Finally, (Fin) follows immediately from (Compactness*). ■

Remark. The set of conditions for Bolzano consequences has a more general significance as also axiomatising relations of functional dependency in database theory. Indeed, the four conditions above are presented in Ullman (1988, p. 384) as a characterisation of functional dependency; they are equivalent to the Armstrong axioms (Armstrong, 1974). We shall come back to this fact at the end of Chapter Five.

Note that inclusion follows immediately from overlap and that weakening and isotonicity are equivalent given the other two conditions. Only one of the conditions characterising Bolzano consequences is not satisfied by Scott consequences: this is the condition of transitivity. Of course, with B restricted to a singleton set, transitivity emerges as a special case of cut and, hence, holds, thus restricted, for Scott consequences; they do not, however, satisfy the fully general transitivity condition. From transitivity and inclusion it follows that Bolzano consequence, while monotonic on the left, is *antitonic on the right*:

$$B' \subseteq B \ \& \ A \Vdash B \implies A \Vdash B'$$

Theories. A set of sentences closed under a consequence operation Cn is a *theory* (according to Cn). Thus, a set A of sentences is a theory just in case

$$\text{Cn}(A) \subseteq A$$

(the converse inclusion holds by condition (Refl) on Cn) or, equivalently,

$$A \vdash \alpha \implies \alpha \in A$$

(the converse holds by the overlap condition on \vdash).

The idealisations and abstractions leading to this definition of a theory are summarised in the figure below.

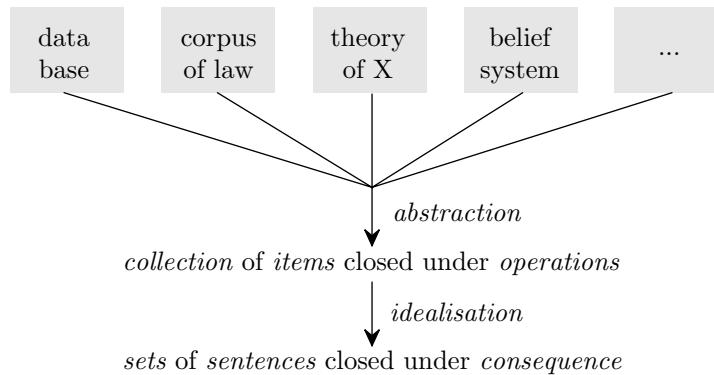


Fig. 2.1

The logical concept of a theory arises as an abstraction and idealisation of various kinds of theories in the ordinary sense. Whatever else theories may be, they certainly have *parts*. They are therefore collections of items. Moreover, not any odd bag of parts is given the title ‘theory’. To be a theory a certain *cohesiveness* is required. A fairly minimal way of cashing in the requirement of cohesiveness is to say that there is some relation such that if an item is in the theory, and if that item bears the relation to some other item, then that other item is in the theory as well. Theories satisfy therefore certain closure requirements.

In the definition of a logical theory, this innocuous abstraction is given a substantial interpretation. Collections are represented as sets, items as sentences, and the closure requirement is given substance by requiring closure under logical consequence.

We have boiled down concrete, everyday theories to a rather austere abstraction. The price to be paid for this abstraction is that in the process we have deliberately lost sight of certain features of concrete theories that no doubt are interesting from one perspective or another. We deliberately disregard, for example, that

- theories are usually united by a subject matter;
- some parts of some theory are closer to direct observation than others;
- some theories (such as mathematical theories) do not rely on observation at all;
- some theories carry a normative force (moral or legal codes) others are purely descriptive;
- many descriptive theories seek explanations;
- many theories are put forward by people who are subject to forces that do not always promote the search for truth.

We have gained, however, in generality, since we have abstracted to the most general properties of theories: what we shall say henceforth will apply to all theories, in the widest sense. But the gain will only be worthwhile, if we find interesting problems that arise and can fruitfully be investigated at the level of abstraction proposed. To such problems we turn now.

2.2 Three kinds of theory change

Theories change in response to certain triggers of change. A trigger can be the sudden availability of new and relevant information which may strengthen, complete or undermine the theory. “Irrational” factors, such as historical, social, or psychological pressure may lead to changing a theory.




Fig. 2.2

An investigation of the triggers of theory change is not part of the theory to be developed here. Any view about the triggers of change is compatible with the logic of theory change.

Also, certain “large scale” changes fall outside the scope of the present theory. Thus, meaning shifts in the terms of a theory or so-called paradigm replacements (in the sense of Kuhn (1970)) are not among the topics to be dealt with. Instead, the AGM theory is about *very small changes*, that is, sentence-by-sentence adjustments. (Later, we shall extend the theory so as to encompass changes by sets of sentences.) Thus, a change operation, in the sense of AGM, takes a theory T and a sentence α to deliver a successor to T . Three directions of change exhaust the range of possibilities:

- *expansions*: the result of expanding a theory T by a sentence α is a larger theory;
- *contractions*: the result of contracting T by α is a smaller theory;
- *revisions*: the result of revising T by α is a theory that neither extends nor is part of the original theory.

These are only rough characterisations. We shall see shortly that there are degenerate cases of the three kinds of changes which do not enlarge, shrink respectively shift incomparably the original theory.

Let us consider an example theory, $T = \text{Cn}(p, p \rightarrow q)$, and ask how it may be expanded, contracted or revised.

Expansion. Suppose we obtain new information r which we should like to incorporate into our theory T . If r is consistent with T , we may just add r to T and close the result under logical consequence.




Fig. 2.3: Expansion

Thus the expansion of a theory may simply be defined in terms of the consequence operation underlying the notion of a theory.

DEFINITION 2.6 The *expansion* of a set A by a sentence α :

$$A + \alpha := \text{Cn}(A \cup \{\alpha\})$$

Contraction. Suppose that we learn that q is not the case. We now need to remove q from T . But q can be removed from T in many ways. We can either remove p or remove $p \rightarrow q$ or remove even both.




Fig. 2.4: Contraction — which way?

(Note that—depending on the number of atomic sentences in a language—there may be many more such subtheories that just the three indicated in this or the next figure.)

Obviously the information we have about T —i.e. $T = \text{Cn}(\{p, p \rightarrow q\})$, where Cn is some consequence operation—does not suffice to nominate non-arbitrarily one of the many subtheories of T not containing q as the contraction of T by q .

Revision. A similar situation arises when we try to revise T . Suppose we obtain new information that $\neg q$ is the case. We should like to adjust T so as to consistently include $\neg q$. As in the case of contraction, this requires an incision into the theory so as to effectively remove q from T .




Fig. 2.5: Revision — which way?

Again, we are faced with a choice which cannot be resolved without taking into account further, as yet unspecified properties of T .

Whereas expansions may be defined solely in terms of logical consequence, further parameters are needed in order to resolve the choice situation faced in our simple example above. There is one constraint on resolving the choice situation which derives from the fact that information is precious and should not be discarded without necessity. This is the

Maxim of Minimal Mutilation:

Keep incisions into theories as small as possible!

At the centre of the AGM theory are a number of approaches to giving formal substance to the maxim. As it turns out, all of these approaches are equivalent. These representation results lend an impressive stability to the AGM theory.

We shall briefly present three approaches in the remainder of this chapter. One proceeds in terms of sets of postulates (the AGM postulates) defining contraction and revision operations. The other two approaches provide direct recipes for how to solve the contraction, resp. revision problem for any given theory and sentence. The one of these approaches assumes the availability of a selection function (possibly generated from a preference relation) on the family of subtheories of the given theory. The other approach assumes that sentences in the theory are more “entrenched” than others, an assumption which may be captured by postulating an (entrenchment) ordering of all sentences in the theory.

2.3 The AGM postulates

Let us assume that a sentential language (with its formulae collected in Fml) and a consequence operation, Cn , on that language have been fixed. Then we have a well-defined space of theories (according to Cn) in that language. The set of theories is bounded by $Cn(\emptyset)$ (the weakest theory, usually the set of theorems of the underlying logic) and Fml (the strongest theory, usually identical with $Cn(\perp)$)

Contractions and revisions are moves in the space of theories governed by two families of postulates as detailed in the table on the next page.

For a detailed motivation of the postulates, the reader may consult Gärdenfors (1988). Here it will suffice to give “plausibility-enhancing” rephrasals of the conditions.

The two *closure* conditions require contractions and revisions to stay within the space of theories. Moving outside this space incurs a serious and unnecessary quality deterioration.

The *congruence* conditions require that the result of revisions or contraction should not depend on syntactic properties of the sentences to be revised or contracted by: only their logical content should count.

If a sentence follows from the empty set, then it follows from any set (by the monotonicity of Cn) and is thus part of every theory. Such sentences cannot be removed. Otherwise contractions are *successful*: a sentence to be contracted by will not be in the contracted theory. For revisions, there are two aspects to success. First, the sentence to be consistently added must

Basic Postulates for Contractions

| | |
|---|-----------------|
| $T - \alpha = \text{Cn}(T - \alpha)$ | closure (C1) |
| $\alpha \in T - \alpha \implies \vdash \alpha$ | success (C2) |
| $T - \alpha \subseteq T$ | inclusion (C3) |
| $\alpha \notin T \implies T \subseteq T - \alpha$ | vacuity (C4) |
| $\alpha \dashv\vdash \beta \implies T - \alpha = T - \beta$ | congruence (C5) |
| $T \subseteq \text{Cn}((T - \alpha) \cup \{\alpha\})$ | recovery (C6) |

Supplementary Postulates

| | |
|---|-------------------|
| $(T - \alpha) \cap (T - \beta) \subseteq T - (\alpha \wedge \beta)$ | intersection (C7) |
| $\alpha \notin T - (\alpha \wedge \beta) \implies T - (\alpha \wedge \beta) \subseteq T - \alpha$ | conjunction (C8) |

Basic Postulates for Revisions

| | |
|---|-------------------|
| $T * \alpha = \text{Cn}(T * \alpha)$ | closure (R1) |
| $\alpha \in T * \alpha$ | success (R2) |
| $T * \alpha \subseteq T + \alpha$ | inclusion (R3) |
| $\neg \alpha \notin T \implies T + \alpha \subseteq T * \alpha$ | preservation (R4) |
| $\neg \alpha \in T * \alpha \implies \vdash \neg \alpha$ | consistency (R5) |
| $\alpha \dashv\vdash \beta \implies T * \alpha = T * \beta$ | congruence (R6) |

Supplementary Postulates

| | |
|--|-------------------------------|
| $T * (\alpha \wedge \beta) \subseteq T * \alpha + \beta$ | conjunctive inclusion (R7) |
| $\neg \beta \notin T * \alpha \implies T * \alpha + \beta \subseteq T * (\alpha \wedge \beta)$ | conjunctive preservation (R8) |

Table: The AGM postulates

be in the revised theory. Second, this is the requirement of *consistency*, the resulting theory must be consistent — again, logic permitting.

Contractions usually remove sentences from a theory; in any case, so the

inclusion condition, they do not enlarge a theory. However, if a sentence to be retracted is not part of the theory, then the contraction operation is *vacuous*.

Similarly with revisions: here the *inclusion* condition gives expression to the idea that revision has both an addition and a subtraction component. But sometimes—when a sentence can consistently be added—, the subtraction component need not be exercised: *preservation*.

A contraction is successful, if the sentences α to be contracted by has been removed. However, if α follows from the empty set (i.e. is a logical truth), then it follows from *any* set (by the monotonicity of consequence). In particular, α follows in this case from any theory contracted by α . Hence, any attempt at contracting α is doomed to fail—in this sense, a contraction by α will be *vacuous*.

The vacuity condition for contractions gives already some expression to the maxim of minimal mutilation: no (proper) contraction without necessity. But the maxim really resides in the *recovery* condition. Recovery requires a contraction of T by α to be such that it can be revoked. Enough must be left in the contracted theory, $T - \alpha$, so as to recover the original theory T once we add α to $T - \alpha$.

Intersection expresses the idea that whatever survives removal of both α and β must also survive removal of $\alpha \wedge \beta$. For to remove $\alpha \wedge \beta$, it suffices to remove α or β —and if we cannot make up our minds, we shall at most remove both. Conjunction says: if we remove α (along with $\alpha \wedge \beta$), then $T - \alpha \wedge \beta$ can be no stronger than $T - \alpha$.

The corresponding supplementary postulates for revisions are motivated similarly. It will here suffice to note that *conjunctive inclusion* generalises inclusion (let β be α and use success) and that *conjunctive preservation* would generalise preservation (let α be \top), if we added the condition that $T * \top = T$ for consistent T . (That latter condition is, however, a special case of preservation.)

The two pairs of supplementary postulates are mentioned only for the sake of completely presenting the AGM postulates; they will play next to no role in what follows.

Whenever one is confronted with a set of postulates to characterise some problematic notion, the question naturally arises: Why these and not other postulates? Eventually we shall see that the chosen sets of postulates provide a *stable* characterisation of contractions and revisions. That is to say, there are a number independently motivated approaches (some of them have already been outlined above) which turn out to agree in their pronouncing a mapping from theories *cum* formulae into theories a contraction, respec-

tively a revision operation.

But even without this external evidence (from independently motivated but converging approaches) there are ways of arguing that the postulates are not badly chosen. In the next subsection we show how the two sets of conditions on contractions and revisions mutually and fully support each other. We close the present subsection with a brief discussion of two pairs of *prima facie* plausible conditions which, on reflection, we have done well not to include among the postulates. The discussion may also serve to further sharpen intuitions about how theories ought to be contracted and revised.

First, consider

$$\beta \in T \ \& \ \beta \notin T - \alpha \implies \alpha \in T - \alpha + \beta \quad (\text{fullness})$$

The condition of fullness seems attractive because it looks like a good candidate for implementing the maxim of minimal mutilation. For fullness says: remove β only if it implies α in the context of $T - \alpha$. But fullness is equivalent to

$$\beta, \gamma \in T \ \& \ \beta \vee \gamma \in T - \alpha \implies \beta \in T - \alpha \ \text{or} \ \gamma \in T - \alpha \quad (\text{primeness})$$

(AGM (1985), Observations 6.1 and 6.2). Primeness, however, is surely an undesirable property for contractions to have. We do want to allow disjunctions in a theory without either disjunct being in the theory. Often we believe with good reason that either β or γ is the case without having sufficient evidence at hand to base this belief on either believing β or on believing γ . In particular, we may start off with beliefs in both β and γ (and, hence, in $\beta \vee \gamma$) and then acquire new evidence to the effect that these beliefs were unwarranted. But the new evidence may not be strong enough to undermine the belief in $\beta \vee \gamma$. In this situation the belief in $\beta \vee \gamma$ should be retained (precisely in the interest of minimal mutilation!) even though both disjuncts had to be retracted—contrary to what primeness would require:




Fig. 2.6: Counterexample to primeness

Second, consider two monotonicity conditions for revisions:

$$T \subseteq T' \implies T * \alpha \subseteq T' * \alpha \quad (\text{monotonicity}^*)$$

$$\alpha \vdash \beta \implies T * \beta \subseteq T * \alpha \quad (\text{cut}^*)$$

The first of these conditions becomes tempting if the analogy between contraction and (set-) subtraction is carried too far. For, one could argue as follows. If $T \subseteq T'$ and we need to remove some set A of sentences from T' in order to pave the way to consistently adding α to T' , then removing A or some subset of A from T will suffice to open T towards adding α . So making T consistent with α results in a set which will be included in the result of making T' consistent with α , and this inclusion will be preserved as we enlarge each of these two sets with α . This train of thought relies on a monotonicity condition for contractions:

$$T \subseteq T' \implies T - \alpha \subseteq T' - \alpha \quad (\text{monotonicity}^-)$$

But this condition may easily fail. Here is a small counterexample: Let T contain $p \rightarrow q$ but not q and let T' be $T + p$. Now suppose we are to contract by q . Then the conditions of inclusion and vacuity force $T - q = T$ while something has to be removed from T' in the transition to $T' - q$. If we choose to remove $p \rightarrow q$ rather than p —and why shouldn't we thus choose?—, then monotonicity for contractions will be violated.

The condition of cut for revisions is open to similar counterexamples. Instead of providing an abstract counterexample, a concrete one will make the failure of the condition more vivid. Let α stand for *this match is struck* and let β stand for *this match is struck* \wedge *this match is wet*. So $\alpha \vdash \beta$ and yet there are plausible T for which $T * \alpha$ but not $T * \beta$ will contain the sentence *this match will light*. A similar counterexample, as sketched in the figure above, illustrates how monotonicity for revisions may fail.




Fig. 2.7: Counterexample to revision-monotonicity

Mutual support: The Levi and the Gärdenfors identities. It will have occurred to the reader that contractions and revisions may be interdefinable. If a sentence α is to be consistently added to a theory T , then one should first make the theory consistent with α and then expand by α . That is to say, to revise T by α should amount to first contracting T by $\neg\alpha$ and then expanding the result by α . This idea is given concise expression in the

$$\textit{Levi Identity: } T * \alpha = T - \neg\alpha + \alpha \quad (\text{LI})$$

(Here and in the sequel we assume that brackets associate to the left. Thus, $T - \neg\alpha + \alpha$ is to be read as $(T - \neg\alpha) + \alpha$ —i.e. $\text{Cn}((T - \neg\alpha) \cup \{\alpha\})$ —rather than as $T - (\neg\alpha + \alpha)$ —i.e. $T - \text{Cn}(\{\neg\alpha, \alpha\})$. The latter would be a case of so-called multiple contraction which will be treated in Chapter Three.) The Levi identity was first explicitly proposed by Isaac Levi in (1977) but it is implicit in many of his earlier writings.

Note that reversing the Levi identity will not result in satisfactory contractions. For in the principal case where T is inconsistent with the sentence α to be revised by, the classical inference *ex falso quodlibet* will explode $T + \alpha$ to the trivial theory Fml , consisting of all sentences of the language. It is not clear how contracting Fml by $\neg\alpha$ can lead us anywhere close to where we intuitively expect $T * \alpha$ to be. Hansson (1992c) reverses the Levi identity in a rather special sense. He proposes to first add α to some finite representation (“base”) S of T without closing under consequence and then to retract $\neg\alpha$ from $S \cup \{\alpha\}$.

The definition of contractions from revisions is less transparent but, on reflection, equally plausible. Suppose α is to be retracted from T . The revision of T by $\neg\alpha$ will (usually) be a consistent set containing $\neg\alpha$. Hence, α will not be in $T * \neg\alpha$. But $T * \alpha$ may be bigger than the target set $T - \alpha$: for one, it will contain $\neg\alpha$. To obtain a subset of T that does not contain α we intersect $T * \neg\alpha$ with the original theory T . Thus results the

$$\textit{Gärdenfors Identity: } T - \alpha = T * \neg\alpha \cap T \quad (\text{GI})$$

It turns out that through the Levi and the Gärdenfors identities the postulates for contractions on the one hand and the postulates for revisions on the other hand completely support each other.

THEOREM 2.7 (*Gärdenfors 1982.*)

1. If the function $- : \wp(Fml) \times Fml \longrightarrow \wp(Fml)$ satisfies the AGM postulates (C1-5, 7, 8) for contractions, then the function $*$ defined from $-$ by (LI) satisfies the AGM postulates (R1-8) for revisions.
2. If the function $* : \wp(Fml) \times Fml \longrightarrow \wp(Fml)$ satisfies the AGM postulates (R1-8) for revisions, then the function $-$ defined from $*$ by (GI) satisfies the AGM postulates (C1-8) for contractions.

Note that the contraction postulate (C6) of recovery, $T \subseteq T - \alpha + \alpha$, is not needed for the derivation of the revision postulated from the contraction postulates via the Levi identity. However, recovery does follow from the revision postulates via the Gärdenfors identity. Since the peculiar status of the recovery postulate is not our present concern, we refer the interested reader to Makinson (1987) for further observations and discussion.

The last theorem greatly simplifies the task for a theory of theory change. We may focus on either contraction or revision and then transfer results accordingly. In the sequel the notion of contraction will be taken as primitive and revision will be treated as defined. Though, formally, we may with equal justification approach theory change by focusing on revisions, the notion of a contraction seems conceptually simpler and the Levi identity more transparent.

2.4 Constructions of contractions

In this section two approaches to constructing the contraction of a given theory by a given sentence will be briefly rehearsed. The main purpose is to introduce and define some concepts which will then be used for the general theory of contractions to be presented in the next chapter.

Besides modellings in terms of so-called partial meet constructions or epistemic entrenchment relations there are at least five more approaches to modelling contractions (or revisions). Under suitable and natural assumptions all these approaches are essentially equivalent in that they pick out the same class of contraction functions. Some of the relevant representation results are surveyed in Rott (1991, 1992b).

Since none of the five alternative approaches will play any explicit role in what is to follow below, we may leave the subject with brief descriptions and some pointers to the literature.

The method of *safe contraction* (Alchourrón and Makinson 1982, Fuhrmann 1991) is, in a way, a mirror image of the partial meet approach to be introduced below. Instead of maximising on the property of not entailing a particular sentence (as in the partial meet approach), it prunes the minimal subsets of a theory that entail the sentence to be retracted.

The *spheres models* of Grove (1988) generalise Lewis' (1973) semantics for counterfactual conditionals in terms of systems of sphere of possible worlds. In Section 3.12 we shall connect our solution to the subtraction problem to possible worlds models in the spirit of Lewis and Grove.

The *minimal models* approach of Katsuno and Mendelzon (1991,1992) is inspired by Shoham's (1988) modelling of reasoning from deductively insufficient premisses (nonmonotonic, or defeasible inference); see Chapter Four.

Among quantitative approaches we mention Shackle measures (Shackle 1961). Epistemic entrenchment relations, which will be described in more detail in a moment, can be shown to be the qualitative part of Shackle measures. This connection is much employed—independently of Shackle—in the work of Dubois and Prade (1992). Spohn (1988) follows another quantitative approach. He uses conditionalisations of *ordinal conditional functions* to represent revisions of theories.

Two approaches to modelling change operations have proved to be particularly flexible and informative: the one proceeds in terms of a relation of *epistemic entrenchment*, the other uses a selection or choice function and intersects selected sets (*partial meet*). The latter approach will be the principal modelling technique in this essay. But before we present this approach, the reader may want to acquaint himself with the principal alternative.

Epistemic entrenchment. (This section may be skipped without impairing understanding of what follows.) While all sentences in a theory must count as fully accepted, some are more accepted than others. For example, some sentences are more central to the concerns of the theory in question than others; or some sentences may be better supported than others; or for some sentences their possibility of falsehood is more remote than for certain others. There are many sources for ranking the sentences in a theory. Without going into a detailed analysis of such sources it suffices here to note that some such sources lead to an ordering of sentences that will constrain the way in which a theory may change, if it has to be changed.

The approach to modelling contraction in terms of an ordering of “epistemic entrenchment” is based on the assumption that each theory may be equipped with a way of ranking its sentences (or all sentences of the underlying language) such that when it comes to choosing between candidates for removal, the least entrenched ones ought to be given up.

DEFINITION 2.8 For each theory T , a relation $\leq_T \subseteq \text{Fml}^2$ is an *epistemic entrenchment (EE) ordering* (for T) if it satisfies the following conditions (subscript T omitted):

$$\alpha \leq \beta \ \& \ \beta \leq \gamma \implies \alpha \leq \gamma \quad \text{transitivity (EE1)}$$

$$\begin{aligned}
\alpha \vdash \beta &\implies \alpha \leq \beta && \text{dominance (EE2)} \\
\alpha \leq \alpha \wedge \beta \text{ or } \beta \leq \alpha \wedge \beta &&& \text{conjunctiveness (EE3)} \\
\alpha \notin T \iff \alpha \leq \beta, \text{ if } \perp \notin T &&& \text{minimality (EE4)} \\
\alpha \leq \top &&& \text{maximality (EE5)}
\end{aligned}$$

Observe that \leq is connected,

$$\alpha \leq \beta \text{ or } \beta \leq \alpha$$

(by EE1–3), that

$$T = \{\alpha: \perp <_T \alpha\}, \text{ if } T \text{ is consistent}$$

(by EE2 and 4), and that the maximal elements under \leq are exactly the logical truths, i.e. $\text{Cn}(\emptyset)$ (by EE2 and 5). (The strict relation $<$ is defined as usual: $\alpha < \beta$ is short for: $\alpha \leq \beta$ but not conversely, $\beta \leq \alpha$.)

One obtains *contraction functions from EE-relations* by letting $T - \alpha$ (in the principal case that $\alpha \notin \text{Cn}(\emptyset)$) consist of exactly those members of T that are more entrenched than α :

$$\beta \in T - \alpha \iff \beta \in T \text{ and } \begin{cases} \alpha <_T \alpha \vee \beta & \text{(the principal case), or} \\ \alpha \notin T & \text{(the vacuous case), or} \\ \top \leq_T \alpha & \text{(so } \alpha \in \text{Cn}(\emptyset)) \end{cases} \quad (\text{OC})$$

(For technical reasons—to do with the verification of recovery—one needs the condition $\alpha < \alpha \vee \beta$ rather than the expected $\alpha < \beta$.)

One obtains *EE-relations from contraction functions* by letting $\alpha \leq_T \beta$ hold just in case one decides against α —but possibly also against β —when either α or β needs to be removed. Thus,

$$\alpha \leq_T \beta \iff \alpha \notin T - \alpha \wedge \beta \quad (\text{CO})$$

Gärdenfors and Makinson (1988) prove the following:

THEOREM 2.9

1. If \leq_T is an EE-relation for a theory T , then the function $-: \{T\} \times \text{Fml} \longrightarrow \wp(\text{Fml})$, defined according to (OC), satisfies (C1–8) (i.e. $-$ is a contraction function over T).
2. if $-$ is a contraction function over a theory T , satisfying (C1–8), then the relation \leq_T , defined according to (CO), satisfies (EE1–5) (i.e. \leq_T is an EE-relation for T).

Let \mathbf{E}_T and \mathbf{C}_T denote the classes of EE-relations for T and of contraction functions over T respectively. The last theorem says that there

exists a mapping c which, according to (OC), takes us from \mathbf{E}_T into \mathbf{C}_T and another mapping e which, according to (CO), maps \mathbf{C}_T into \mathbf{E}_T (for all theories T). The result can be strengthened by observing (as in Gärdenfors and Makinson 1988) that c and e establish a bijection between \mathbf{E}_T and \mathbf{C}_T : if we apply c to a relation $\leq \in \mathbf{E}_T$, then applying e to $c(\leq) \in \mathbf{C}_T$ will return us to the original relation \leq . The same holds, *mutatis mutandis* for e ; hence,

$$ce = ec = \textit{Identity}$$

Partial meet. In the AGM tradition, the most important definition of singleton contraction is that of *partial meet* (p.m.) contraction; see AGM (1985). The basic idea is that one should try to lose as little information as possible when contracting a belief set. Suppose we are to contract a theory T by some sentence α . As a first approximation towards contracting without incurring loss of information beyond necessity, we may restrict attention to the *maximal* subsets of T that do not entail α . Call such subsets of T *remainders* and let $T \perp \alpha$ be the set of remainders (of T after removing α).

In all non-degenerate cases there are many such remainders; in fact, there are too many remainders to let their intersection (so-called full meet contraction) be a viable candidate for the contraction of T by α . On the other hand, picking an arbitrary remainder brings in an element of gambling where rational choice is asked for. Besides, remainders are in a way “too large” to qualify as candidates for the contracted theory, as in so-called maxichoice contraction; see AGM (1985) for some negative results. Thus it appears that the consequence operation on its own—or, rather, in conjunction with familiar set-operations—does not suffice to determine that successor to a given theory that deserves the title “contraction” (by some given sentence).

At this point we simply pad our logical apparatus with a brute but natural assumption: that among a collection of alternative remainders we can somehow pick those that are, in some sense, the most preferred ones in that collection. Note that it is not assumed that the choice can always be narrowed down to uniqueness: there may be more than one most valuable remainder. Given that we have revealed our preferences by choosing a set of remainders we define $T - \alpha$ to be the set of sentences that are common to all preferred remainders, i.e.

$$T - \alpha = \bigcap s(T \perp \alpha)$$

where s is a mapping (a selection function) from a non-empty class of theories (the remainders) into a non-empty subset of that class (the preferred

remainders). We need to take care of the special case when $T \perp \alpha$ is empty. This happens just in case α cannot be removed because it is a logical truth. In that case we simply put $s(T \perp \alpha) = \{T\}$ whence $T - \alpha = T$.

Partial meet contractions may be defined in three steps. First, the definition of remainder sets:

DEFINITION 2.10 (Remainders) For each theory T and sentence α , $X \in T \perp \alpha$ if and only if

- (a) $X \subseteq T$,
- (b) $X \not\vdash \alpha$, and
- (c) $\forall Y : X \subset Y \subseteq T \implies Y \vdash \alpha$.

Sets in $T \perp \alpha$ will be called *remainders* (of T after removing α).

Note that remainders are closed under consequence:

LEMMA 2.11 *If A is a remainder, then $\text{Cn}(A) = A$.*

Next we introduce a function which chooses “preferred” remainders.

DEFINITION 2.12 A *selection function* (for a theory T) is any function

$$s_T: \wp(\wp(T)) \longrightarrow \wp(\wp(T))$$

such that $\emptyset \subset s_T(\mathcal{X}) \subseteq \mathcal{X}$ for all $\mathcal{X} \neq \emptyset$, and $s_T(\mathcal{X}) = \{T\}$ otherwise. (In the sequel subscripts to s will be omitted wherever convenient.)

Now we are in a position to define partial meet contractions.

DEFINITION 2.13 A function $-: \wp(Fml) \times Fml \longrightarrow \wp(Fml)$ is a *partial meet (pm) contraction* over a theory T if and only if there exists a selection function s_T such that for all $\alpha \in Fml$,

$$T - \alpha = \bigcap s(T \perp \alpha)$$

THEOREM 2.14 (AGM, 1985, Observation 2.5.) *For each theory T , a function $-: \wp(Fml) \times Fml \longrightarrow \wp(Fml)$ is a pm contraction over T if and only if it satisfies the basic contraction postulates (C1-6) for T .*

PROOF. The argument proceeds as follows. From left to right the assertion is verified by showing that $T - \alpha = \bigcap s(T \perp \alpha)$ satisfies the postulates. For the converse direction one defines a “canonical” selection function (for each theory T):

$$s(T \perp \alpha) = \begin{cases} \{T' \in T \perp \alpha : T - \alpha \subseteq T'\} & \text{if } T \perp \alpha \neq \emptyset \\ \{T\} & \text{otherwise} \end{cases}$$

One needs to show that

- (1) s is well-defined, i.e. $T \perp \alpha = T \perp \beta \implies s(T \perp \alpha) = s(T \perp \beta)$;
- (2) $s(T \perp \alpha) = \{T\}$ if $T \perp \alpha = \emptyset$, which is immediate from the definition;
- (3) $s(T \perp \alpha) \subseteq T \perp \alpha$ if $T \perp \alpha \neq \emptyset$, which is likewise immediate from the definition;
- (4) $s(T \perp \alpha) \neq \emptyset$; and
- (5) $T - \alpha = \bigcap s(T \perp \alpha)$.

For (1) one uses extensionality. The other postulates enter the verification of (4) and (5). Recovery is needed for the latter. \blacksquare

It can also be shown (AGM 1985, Observations 6.2 and 6.3) that the fullness condition, discussed on page 23 above, corresponds to the assumption that each selection function always picks out a unique remainder (*maxi-choice* contraction) and that the cut condition for contractions,

$$\alpha \vdash \beta \implies T - \alpha \subseteq T - \beta, \text{ for all } \alpha \in T \quad (\text{cut}^-)$$

corresponds to the assumption that the selection function is trivial in that it amounts to the identity function: $s(\mathcal{F}) = \mathcal{F}$, for each nonempty family \mathcal{F} of remainder sets (*full meet* contraction). A condition equivalent to cut for contractions which is sometimes used to characterise full meet contraction is

$$T - \alpha \wedge \beta = T - \alpha \cap T - \beta \quad \text{full intersection}$$

We have already (p. 23f.) adduced a counterexample to fullness. It is an undesirable property for a contraction operation; hence, contractions should not generally be maxichoice. Continuing the discussion of cut principles (p. 23) we report a result of Alchourrón and Makinson (1982, Observations 2.1 and 2.2) concerning full meet contractions, i.e. contractions satisfying, apart from the basic postulates, the condition of cut for contractions (equivalently: full intersection). The result shows that full meet contractions are “too small” and that they trivialise associated revision operations.

OBSERVATION 2.15 *If $-$ is a full meet contraction operation, then for every $\alpha \in T$,*

1. $T - \alpha = T \cap \text{Cn}(\neg\alpha)$;
2. $T - \neg\alpha + \alpha = \text{Cn}(\alpha)$.

It is essential for this result that T is a theory, i.e. closed under consequence.

To model the supplementary postulates of intersection (C7) and conjunction (C8) one assumes that for each theory T there is a relation \leq_T on the family $\wp(T)$ of subsets of T . Intuitively, \leq_T orders the parts of T according to their relative importance for the theory as a whole; one may

think of Quine's web of belief (Quine and Ullian 1978). Though this interpretation immediately confers some basic properties, such as transitivity, onto the relation, no appeal to such properties is needed to model (C7).

Let us say that a selection function s_T is *generated* from a relation \leq_T on $\wp(T)$ just in case for each family \mathcal{F} of remainder sets (of T after removing some sentence), $s_T(\mathcal{F})$ coincides with the \leq_T -maximal elements of \mathcal{F} :

$$s_T(\mathcal{F}) = \{X \in \mathcal{F} : Y \leq_T X \text{ for all } Y \in \mathcal{F}\}$$

Thus, if $X \leq Y$, then Y is “(weakly) preferred to” or “at least as good as” X .

DEFINITION 2.16 A function $-: \wp(Fml) \times Fml \rightarrow \wp(Fml)$ is a (*transitively*) *relational pm contraction* over a theory T if and only if there exists a selection function s_T such that for all $\alpha \in Fml$,

$$T - \alpha = \bigcap s(T \perp \alpha)$$

where s_T is generated from a (transitive) relation $\leq_T \subseteq \wp(T)^2$.

Every relational pm contraction satisfies the condition (C7). If the generating relation is transitive, then (C8) will also be satisfied (AGM 1985, Observation 4.3). Moreover, the conditions (C1–8) completely characterise the class of transitively relational pm contractions. However, the larger class of simply relational pm contractions are *not* characterised by (C1–7). That is, contrary to what one may have expected, (C8) does not pick out the transitivity of the generating relation. An exact characterisation of relational pm contractions is one of the most prominent open problems in the AGM theory (cf. Rott 1993). Thus, for the time being, we must rest content with recording the following result.

THEOREM 2.17 (AGM 1985, Observation 4.4.) *A function $-: \wp(Fml) \times \wp(Fml) \rightarrow \wp(Fml)$ is a transitively relational pm contraction over T if and only if it satisfies the contraction postulates (C1–8) for T .*

PROOF. For each T , the canonical relation $\leq_T \subseteq (\wp(T))^2$ may be defined as follows: $X \leq_T Y$ if and only if both X and Y are members of some remainder sets of T and either

1. the trivial case $X = T = Y$ holds, or
2. X and Y are members of some remainder sets of T and
 - $T - \alpha \subseteq X$, for some $\alpha \in T$ and
 - for any α and $X, Y \in T \perp \alpha$: if $T - \alpha \subseteq Y$, then $T - \alpha \subseteq X$. ■

General Contraction

In this chapter the AGM theory of belief contraction will be completely generalised to contractions of arbitrary sets of sentences by arbitrary sets of sentences. The resulting theory will encompass, as special cases, theories of

- how to contract logically closed sets by single sentences (singleton contraction, AGM (1985)),
- how to contract logically closed sets by sets of sentences (multiple contraction, Fuhrmann and Hansson (1994)), and
- how to contract arbitrary sets of sentences by single sentences (base contraction, Fuhrmann (1988, 1991), Hansson (1989), Nebel (1991)).

The items to be contracted will be called *belief sets*. These are simply sets of sentences. In contrast to the usage of the term in Gärdenfors (1988) and some of his earlier writings, belief sets need not but can be closed under a suitable consequence operation. Sometimes we shall stress the fact that a particular belief set is not closed under consequence by calling it *open*. The term ‘belief set’ in the present sense seems to have found general acceptance in the relevant literature. Mainly for this reason it will be used here. The term indicates an important area of application—natural or artificial epistemic states—but is not meant to exclude others such as systems of norms or those indicated in Fuhrmann (1995, 199+).

There are two major variants of the general notion of contraction. In the *package contraction* of a set A by a set B all sentences in B must be removed from A , whereas in *choice contraction* it is sufficient that at least one member of B is removed. The principal modelling tool for both

kinds of contractions will be suitable generalisations of the partial meet construction of AGM (1985). These constructions will be matched by sets of characteristic postulates.

3.1 An Imbalance

All change operations considered in the theory of AGM—expansions, contractions and revisions—operate on a *theory*, say T , and a *single sentence*, say α , to deliver a new, changed, theory, T' , where the change from T to T' is controlled by the sentence α .

The AGM theory of belief change exhibits a curious asymmetry in the arguments that a change operation may take. This asymmetry is most clearly exposed if one looks at the grammar of change operations: these are binary operations that take two sets of sentences, A and B , to return a third set of sentences C :

$$(A, B) \mapsto C$$

Intuitively, the left-hand argument, A , is the item to be changed. The right-hand argument forces—up to uniqueness—a certain “direction of change”: it represents the item to be “changed by”. Asymmetry enters the AGM theory by way of two, possibly three constraints:

Closure: the left argument is closed under consequence;

Singleton: the right argument is a singleton set;

Consistency: the left argument is consistent.

Whereas the first two conditions must clearly be counted among the hallmarks of the AGM theory, the status of the third condition remains doubtful. For example Gärdenfors (1988, 22) appears to make consistency part of the definition of a belief set, the consistency conditions plays no role in AGM (1985). Levi (1991) bases part of his criticism of the AGM theory (mainly as espoused in the writings of Gärdenfors) on the consistency requirement. In Fuhrmann (1991) it is argued that the consistency condition is an undesirable constraint in a theory of belief change and that it need not be taken as part of the AGM approach. In Chapter Four we shall introduce a change operation, “merge”, whose sole purpose it is to salvage as much as seems reasonable from an inconsistent belief set.

In the following two sections the closure and the singleton condition will be discussed in turn.

3.2 Closure

The condition that the objects of change should be closed under consequence marks the main difference between theory contraction and what has come to be called *base contraction* (Hansson (1989), Fuhrmann (1989, 1991)). Note that although the closure condition is now usually associated with the AGM theory, it was not there in the beginning. In the first papers by Alchourrón and Makinson the condition was not required. And in their survey of partial meet contractions (1985) AGM took much care to signal where the condition is actually used.

In base contraction the item to be contracted is allowed not to be closed under consequence. Sometimes this leniency is strengthened to a requirement (as in Hansson (1989)). The observation motivating base contraction is that, as a matter of fact, we never engage directly in theory contraction. Instead we adjust surveyable—and that means at least finite—representations of theories. So all theory contraction should take its route through adjusting a distinguished subset of the theory in question, its base.

The slogan “theory contraction through base contraction” (Fuhrmann 1991) indicates that the move from theories to bases is not meant to entail a move from contraction to set-theoretic subtraction. That would amount to a (near) trivialisation of the notion of a contraction. When in the sequel we say that some, possibly open set A is to be contracted by α or that α should be removed or retracted from A , we mean that A should be adjusted such that α is no longer *derivable* from A . This is of course a much stronger—and much more interesting—requirement than that α is no longer contained in A . Thus, α has been removed from A just in case α is not contained in $\text{Cn}(A)$.

Base contraction does not only feel homelier than the more opaque notion of theory contraction, it also does justice to an important intuitive constraint on theory change. This constraint was introduced as the *filtering condition* in Fuhrmann (1991). The following example, illustrating the need for considering bases and how the filtering condition is supposed to work, is taken from Hansson (1989, pp. 117–8):

“(…) suppose that on a public holiday you are standing in the street in a town that has two hamburger restaurants. Let us consider the subset of your belief set that represents your beliefs about whether or not each of these two restaurants is open.

“When you meet me, eating a hamburger, you draw the conclusion that at least one of the two restaurants is open ($\alpha \vee \beta$). Further, seeing from the distance that one of the two restaurants has its lights on, you believe that this particular restaurant is open (α). This sit-

uation can be represented by the set of beliefs $\{\alpha, \alpha \vee \beta\}$.

“When you have reached the restaurant, however, you find a sign that it is closed all day. (...) You now have to include the negation of α , i.e. $\neg\alpha$, into your belief set. The revision of $\alpha, \alpha \vee \beta$ to include $\neg\alpha$ should still contain $\alpha \vee \beta$, since you still believe that one of the two restaurants is open.

“In contrast, suppose you had not met me or anyone else eating a hamburger. Then your only clue would have been the lights from the restaurant. The original belief system in this case can be represented by the set $\{\alpha\}$. After finding out that this restaurant was closed, the resulting set should not contain $\alpha \vee \beta$, since in this case you have no reason to believe that one of the restaurants is open.”

But $\{\alpha\}$ and $\{\alpha, \alpha \vee \beta\}$ are logically equivalent; they, therefore, determine the same theory. Yet their respective revisions give intuitively different results. If these intuitions are to be respected, revision has to take account of the base from which a theory is generated. Different bases may have different dispositions to change, even if their logical closure is the same.

There is a different way of making essentially the same point (see Fuhrmann (1991), pp. 184–6). The disjunction $\alpha \vee \beta$ is in the theory determined by the base α . But if we revise the *base*—not the theory!—by $\neg\alpha$, the result will be the empty set from which $\alpha \vee \beta$ no longer follows. If, instead, we revise the *theory* determined by $\{\alpha\}$ by $\neg\alpha$, there will be no guarantee that $\alpha \vee \beta$ will have been removed from the theory along with α . In fact, the recovery postulate—expressing a very strong reading of the minimal mutilation maxim—will require $\alpha \vee \beta$ to stay! This violates the firm presystematic datum that the dependency relation generated by the base ought to be respected. The sentences in the base are our building blocks. Everything that follows from the base is in a fairly obvious sense dependent on those blocks. If part of the base goes, so should anything that essentially depended on that part. Hence,

The Filtering Condition:

If β has been removed from a base A in order to bar derivations of α from A , then the contraction of $\text{Cn}(A)$ by α should not contain any sentences which were in $\text{Cn}(A)$ “just because” β was in A .

(For the present purpose we need not here spell out the phrase “just because”; for the requisite details see Fuhrmann (1991, p. 185).)

3.3 Very small changes

The second condition, that all changes are by single sentences, has by and large withstood the tide of time. A first sketch of a theory of change operations that are not generally singular on the right is contained in Fuhrmann (1988). Hansson (1989), Niederée (1991), and Rott (1992) have independently pursued the idea of multiple change operations. In the present chapter we shall probe further into the topic of changes by more than one sentence at a time and eventually present a fully developed theory of multiple change operations.

The term “multiple contraction” was proposed in Fuhrmann (1988) for operations of contraction that allow for simultaneous contraction by more than one sentence. It should be distinguished from “repeated” or “iterated” contraction, i.e., the performance of two or more contractions in a sequence. Multiple contraction, as the notion is to be understood here, is simultaneous contraction by a set of sentences which need not be a singleton set.

It is obvious that base contraction is a more general notion than theory contraction. On the one hand, a theory of base contractions covers theory contractions as special cases, namely as contractions of “closed bases”. On the other hand, there are—very many!—cases of base contractions which cannot be represented in terms of theory contraction. Thus, the theory of base contractions is a *proper* extension of the theory of theory contraction. It is far less obvious, however, whether the notion of a multiple contraction likewise properly extends that of a singleton contraction. To this question we need to turn next.

3.4 Do, need, and should we multiply contract?

It is one thing to generalise the format of contraction operations by allowing not only singleton sets but arbitrary sets as items to be removed. It is quite another thing to attach to this generalisation a thesis to the effect that it is “required” for the purpose of modelling belief change. Such a thesis will indeed be put forward in a moment. But to see what exactly will be claimed here, it will be useful to distinguish three questions:

1. Do multiple contractions occur?
2. Need multiple contractions occur?
3. Should multiple contractions occur?

It will shortly be argued that, as a matter of brute fact, multiple contractions *do* occur in our cognitive practice—though it may require a step

of abstraction to recognize their presence. Thus, a theory of multiple contractions is at least insofar required as actual occurrences of belief change are to be captured.

But we have agreed to develop our theory of contractions at a certain level of abstraction. Thus, it could be argued, that abstraction should lead to discarding the fact that multiple contractions happen to occur.

Multiple contractions may only occur as a matter of convenience. Perhaps they can, at least in principle, always be replaced by singleton contractions and, hence, *need* not occur. If so, we would have a good case for dismissing multiple contractions as theoretically—though perhaps not practically—unimportant.

Below we shall consider a number of candidates for reducing multiple contractions to singleton contractions. However, as we shall see, none of these candidates allow in general for such reduction, though under certain conditions reduction is possible. Given then that multiple contractions are *sui generis*, there is a sense in which they *must* occur: there are no alternatives to achieve their effects.

Even if multiple contractions cannot be reduced to singleton contractions, their rationality may be questioned: *should* they occur? The question will be pursued a little further below. Eventually, however, it is beyond the scope of the present essay and we shall have to leave its further exploration for another occasion.

The ubiquity of multiple contractions. It seems difficult at first to find cases of pure contraction. In most cases, the retraction of a belief is provoked by the acquisition of some other belief which forces the old one out. Contraction thus seems to occur most frequently as part of a more complex operation which involves both the removal and the addition of information. The most clear case of pure contraction that has been discussed in the literature is what may be called ‘contraction for the sake of argument’ or ‘mind-opening contraction’. One may wish to give a belief that α a hearing although it contradicts one’s present state of belief. To open one’s mind for α one should thus retract $\neg\alpha$. (Levi (1980, ch. 3), (1991, ch. 4); Gärdenfors (1988, p. 60); Fuhrmann (1991, Sec. 4).)

Contraction for the sake of an argument will often involve the simultaneous removal of more than one belief. Suppose, for instance, that you wish to give each of two different views—both of which contradict your present standpoint—a fair hearing. You then need to open up your belief state to make it consistent with either of the two views. Representing the two views by the sentences $\neg\alpha$ and $\neg\beta$ respectively, this amounts to removing both α and β from your belief set. The result of this operation should be a belief

set that is open to both $\neg\alpha$ and $\neg\beta$ in that it implies neither α nor β and is otherwise as similar as possible to the original belief set.

Pure contraction, i.e. contraction which does not immediately serve the consistent acquisition of new information, occurs also when one has inadvertently expanded a consistent belief set by some sentence α such that $\neg\alpha$ is in the belief set. As was observed by Levi (1991, Sec. 4.8), in such cases it is often best to restore consistency by retracting both α and $\neg\alpha$. In general there is no reason to contract first by α and then by $\neg\alpha$, or the other way around. The situation is therefore best be viewed as a case of multiple contraction by $\{\alpha, \neg\alpha\}$.

Multiple contraction is applicable in non-epistemic contexts as well. Consider, for instance, the dynamics of normative systems—which, by the way, was one of the main sources for the logic of theory change; see Alchourrón and Makinson (1981). Changes in legal systems often come in large pieces: they typically involve the simultaneous alteration of many parts of a legal code. They are thus typically instances of multiple changes.

It appears that we also retract logically closed parts from theories. That is to say, there are cases of theory contractions by theories. Most theories come in layers. For example, physics is (partly) based on a background of mathematics which in turn is (partly) based on a background of logic. This picture—a heritage from the Viennese Circle—is, perhaps, in some ways too simple but nevertheless essentially correct. The picture does not only apply to the architectonics of science but also models ways in which information is structured in artificial intelligence systems, say, in distributed databases. Consider then a partial order of theories with a maximal element, T , (incorporating all theories below in the order) such that for all theories T_1 and T_2 : if $T_1 \leq T_2$ and $\alpha \rightarrow \beta \in T_1$ and $\alpha \in T_2$, then $\beta \in T_2$. (Probably the first formal study of such structures is contained in Meyer (197+).) There may be occasions when we need to study the effect of removing a whole theory T_i from the partial order—perhaps, in order to replace it by some other theory. (This is like removing a module from a distributive database.) Clearly, this is an instance of multiple contraction.

The necessity of multiple contractions. It is important to distinguish the operation of completely removing a set of sentences, say, $\{\alpha, \beta\}$, from a belief set from several other operations which, at first sight, seem to have similar effects:

1. contracting by $\alpha \vee \beta$,
2. intersecting the results of contracting by α and of contracting by β ,
3. first contracting by α and then by β , or vice versa,
4. contracting by $\alpha \wedge \beta$.

In Section 3.9 we shall investigate possible strategies for reducing contractions by sets to contractions by single sentences in more detail. For now the following informal remarks may serve to indicate that it is far from obvious that such reduction strategies can be successful.

As to the first operation, it is true that in order to remove a disjunction from (the closure of) a belief set one needs to remove both disjuncts. But not conversely: contracting by the set $\{\alpha, \beta\}$ does not require removal of $\alpha \vee \beta$. One may open one's mind to both $\neg\alpha$ and $\neg\beta$ without opening it to $\neg(\alpha \vee \beta)$. For a rather drastic counterexample to this particular reduction thesis let $\beta = \neg\alpha$.

The second operation may result in a belief set which is too small to represent the corresponding multiple contraction. Still, this reduction strategy seems to be the most promising of the four and we shall return to it at the end of Section 3.8.

The third operation would introduce asymmetry where there should be none. For the order of contraction may make a difference (see e.g. Hansson 1992c). We cannot expect in general that the result of first contracting by α and then by β is the same as that of first contracting by β and then by α . But it is part of the very idea of a multiple contraction that all sentences to be retracted are retracted *simultaneously*. Thus, sequential contraction cannot be the same operation as multiple (i.e. simultaneous) contraction.

The temporal order of sequential contraction could be obliterated by intersecting all possible contraction sequences; in the case of a two-elements-set we could intersect the result of first retracting α , then β with the result of first retracting β , then α . But this operation would make the contracted belief set even smaller than the one that results from applying the second strategy above. Given that the second operation cannot serve our purpose (as will be shown below), its symmetrizing mate must fail too.

That the fourth operation is unsuitable as a representation of a contraction by $\{\alpha, \beta\}$ is easy to see: to remove a conjunction it suffices to remove one of the conjuncts. Thus, we may have α staying in a contraction by $\alpha \wedge \beta$ but, of course, α should not stay in a multiple contraction by $\{\alpha, \beta\}$ as we have used the term so far.

Up to now we have tacitly assumed that to multiply retract a set B from A is to remove *all* members of A from B . The remarks of the last paragraph point towards an alternative notion of multiple contraction: instead of removing a certain set completely from a belief set one might be interested in an operation that modifies a belief set such that this set is no longer contained in it. This notion of a multiple contraction is useful whenever a belief set is faced with contravening evidence which, however, is not specific

enough to determine exactly which sentences ought to be retracted in order to accommodate the evidence. This is a situation frequently encountered: the recalcitrant evidence may, on its own, not determine which parts of the belief set ought to be given up—in this sense it may be inconclusive. In the worst case the evidence gives no guidance at all as to which part of one’s beliefs is likely to be false. More usually, however, the evidence will incriminate a proper part of the belief set from which culprits should be picked. In any case we are faced with a choice which can only be made in the context of the whole set of beliefs. That is to say, the theorist has to judge—according to some contextually fixed relevant parameters—which part of the belief set she can best afford to live without. This is precisely the process James describes in the passage quoted at the opening of Chapter One.

(The theory of belief change itself supplies a fine example of this kind of contraction. Gärdenfors’ (1988, Sec. 7.4) impossibility theorem concerning the generation of conditionals in systems of belief revision is commonly taken to require that, in the presence of the Ramsey test, the full set of revision postulates be retracted. But, of course, the retraction ought not to be a complete one but one that involves choice—it is far from clear, however, which choice is to be made.)

We can thus discern two kinds of multiple contraction. According to one kind of contraction all members of a set are retracted: they have to go in a package. There is another kind of contraction where one only needs to ensure that some set is no longer a subset of the belief set in question. For that purpose it suffices to remove some elements of that set from the belief set: one needs to choose which ones. We shall call the first type of multiple contractions *package contraction* and the second type *choice contraction*. It will emerge (Observation 3.17 below) that in the finite case the latter reduces to contraction by the conjunction of all sentences to be retracted.

To see how both package and choice contractions can be useful in the same situation, consider again recovery from inconsistency. Suppose the set $\{\alpha, \neg\alpha\}$ is to be removed from a belief set. If the epistemic agent is only concerned with regaining consistency, then choice contraction is sufficient. If, on the other hand, she is more cautious (or “sceptical” to use a popular AI-term) and wishes to investigate again both the statements that led to contradiction (as Levi (1991) tends to recommend), then a package contraction is called for.

The rationality of multiple contractions. The last question, should multiple contractions occur?, is perhaps the most difficult of the three: given that information is precious, can it ever be rational to engage

in multiple contraction? This is a question I shall not here address. Instead I shall proceed from the assumption that decisions to multiply contract are sometimes taken and that some ways of implementing such decisions are better than others. (Perhaps it is generally irrational to run nuclear power plants. But given that nuclear power plants exist, there are more and less rational ways of operating them.)

Although I shall not here pursue the question as to which contractions ought to be carried out, it is herewith recommended to the reader's attention. The rationality conditions for contractions are studied in depth in many of Levi's writings (e.g. 1980, 1991). According to Levi there are only two situations in which it is rational to contract: opening one's mind for new possibilities and retreating from inconsistency. Levi merges the second ("*need* multiple contractions occur?") with the third question ("*should* they occur?") by proffering the conjecture (in correspondence) that every legitimate contraction, that is, every contraction that serves one of the two purposes, must be (equivalent to) a singleton contraction. This is not a general reduction thesis but one pertaining only to those contractions one *should* engage in, i.e. the legitimate or rational contractions. Evidently, the conjecture turns much on the requisite notion of legitimacy; it is accordingly much more difficult to assess than the general reduction theses considered, and refuted, in Section 3.9.

This concludes, for the time being, our informal motivation and characterisation of contraction by sets. At this stage hardly more can be done than broadly outline the target notions and give some indication that we are engaged in a worthwhile enterprise. Once we are equipped with a corpus of formal definitions and results, we shall return to some of the questions raised in this section.

3.5 Package contraction and choice contraction

As indicated in the last section, multiple contractions come in two varieties: package and choice contraction. The difference is most notable in the two success postulates required for a general theory of contractions

The purpose of the success postulate in the AGM theory,

$$\alpha \notin A - \alpha, \text{ unless } \emptyset \vdash \alpha$$

is to ensure that the sentence to be removed from a theory should no longer be contained in the theory after the contraction has been performed. We need the unless-clause because of the closure condition in the AGM theory. According to that condition, contractions proceed from theories to theories.

Thus, if α is a consequence of the empty set—and, hence, of any set—it cannot be removed.

Although in the general theory closure is not required, logical consequence enters the theory through the intended notion of removal: to remove a sentence from a set A is to perform an incision into A such that the sentence can no longer be derived. Thus, in the general theory too success cannot be unqualified. If a sentence follows from any set, it will in particular follow from contracted sets. Consequences of the empty set (Logical Truths) are immune to effective retraction.

In the sought generalisation of success for singletons the formula α must be replaced by a set B of formulae. Contracting a belief set A by B should “remove” A from B . There is an evident ambiguity here: to remove A from B can mean either

- that none of B is a consequence of A (complete removal), or
- that some of B is not a consequence of A (partial removal).

As argued above, both possibilities are interesting in their own right and need to be investigated.

To completely contract A by B is to remove all elements of B from A . We shall call this kind of removal operation *package* contraction and use the notation

$$A \overset{\forall}{\dashv} B.$$

As a tentative success postulate for package contractions we may try: $A \overset{\forall}{\dashv} B \not\vdash B$ —none of B follow from $A \overset{\forall}{\dashv} B$. (In Fuhrmann (1988) the term “meet” contraction was used instead. The new term avoids confusion with the concept of a partial meet contraction as it occurs in the writings of AGM.)

To contract A so that B is no longer contained in the set of consequences of A it suffices to remove at least one of the elements of B from A . We shall refer to this kind of contraction as a *choice* contraction and use the notation

$$A \overset{\exists}{\dashv} B.$$

A tentative success postulate for choice contraction is: $A \overset{\exists}{\dashv} B \not\vdash B$ —some of B do not follow from $A \overset{\exists}{\dashv} B$.

Sometimes it is convenient to remain ambiguous between package and choice contraction. We shall then speak of (multiple) contraction and use the (ambiguous) notation $A - B$. To aid browsing we shall frequently—but not always!—prefix the name of a postulate with a ‘P’ (for ‘Package’) or a ‘C’ (for ‘Choice’) even when the context leaves little doubt as to which kind of multiple contraction is under consideration.

As noted a moment ago, since removal from A means removal from the logical closure of A , the tentative success conditions need qualification for the case that the underlying logic does not permit removal. If B consists of theorems only, B cannot be choice-removed from any belief set; and if B contains some theorem, then B cannot be completely or package-removed from A . Hence, the proper success conditions are

$$\begin{aligned} A \sphericalangle B \vdash B &\implies \vdash B && \text{P-success} \\ A \cong B \Vdash B &\implies \Vdash B && \text{C-success} \end{aligned}$$

The condition that the set B must contain no theorems for a package contraction by B to be successful, may appear to be a very strong one. To illustrate somewhat drastically, suppose that B is very large but contains a single theorem. Then, for all we know so far, an application of package contraction may not remove any sentence from the original belief set, simply because the condition under which success is guaranteed, does not obtain.

An intuitively more satisfactory variation on the theme of package contractions is a contraction operation which, as it were, tries to do its best: it removes all those sentences in B from a belief set that can be removed. Let \smile stand for such a contraction operation. An appropriate success postulate would be

$$A \smile B \nVdash B \setminus \text{Cn}(\emptyset)$$

where \setminus denotes set-theoretical subtraction. This kind of contraction appears to be well-positioned between choice and package contractions. But it is easily defined in terms of the latter by putting

$$A \smile B := A \sphericalangle (B \setminus \text{Cn}(\emptyset))$$

Some consideration needs to be given to the empty set. What are we to make of contractions such as $\emptyset - B$ or $A - \emptyset$ or even $\emptyset - \emptyset$? We shall legislate the answer “Nothing” by defining contractions between non-empty sets only. To be sure, there are a number of options as to how to extend the theory to these limiting cases. But no presystematic preferences seem to dictate any particular choice and none of them provides additional insight or even a formal smoothing of the theory.

3.6 Basic properties of multiple contraction

In the last section, we have generalised the AGM postulates of closure (C1) and success (C2). In this section, we shall discuss the requisite generalisation of the remaining four basic AGM postulates and consider some closely related properties of contraction functions.

According to the *inclusion* condition, the result of a singleton contraction should be a subset of the original set:

$$A - \alpha \subseteq A$$

The basic intuition behind inclusion is that contraction is “pure” in the sense of not involving the acquisition of any new item of belief. Although it may be maintained that pure contraction in this sense is not common as an isolated phenomenon, it may occur as part of more complex changes in belief. In this sense, pure contraction is at least a useful logical abstraction—like material implication—and an important tool in the analysis of belief change. This holds true for the general concept of a contraction function just as well as for its special case, theory contraction by singletons.

The appropriate generalisation of inclusion seems obvious and it is the same for both choice and package contraction:

$$\begin{array}{ll} A \succ B \subseteq A & \text{P-inclusion} \\ A \exists B \subseteq A & \text{C-inclusion} \end{array}$$

In the next chapter we shall consider an important variation on the theme of removal which differs mainly in the formulation of the inclusion condition.

According to AGM’s *vacuity* postulate, the contraction by something that was not in the original belief set is an idle operation. In other words, if what should be achieved by the contraction has already been achieved, then the operation of contraction is vacuous. If you do not believe that London is the capital of France, then the contraction of your belief set by that belief involves no change at all. In general we should require for single sentences α that

$$\alpha \notin A \implies A = A - \alpha$$

One half of the consequent holds already unconditionally by inclusion; the other half is the vacuity condition. For package contraction, the corresponding principle should come into force only for sets that are completely disjoint from A . Only this case completely anticipates the effect of the proposed contraction. To illustrate this, suppose your beliefs entail that London is the capital of France but not that Berlin is the capital of Germany. Then the

contraction of your belief set by the set containing these two beliefs should be a real change which removes the last-mentioned belief from your set of beliefs. If, on the other hand, you are committed to neither of these beliefs, then the contraction should be vacuous. Thus we have the following vacuity condition for package contraction:

$$A \not\vdash B \implies A \subseteq A \overset{\forall}{\dashv} B \quad \text{P-vacuity}$$

With choice contraction, the intuition behind vacuity gives rise to a different principle. A choice contraction aims at removing at least one element of the set to be retracted. Therefore, for vacuity to come into force, it suffices that there is one element of B that is not entailed by A . Thus,

$$A \not\vdash B \implies A \subseteq A \overset{\exists}{\dashv} B \quad \text{C-vacuity}$$

The most controversial among the six basic Gärdenfors postulates is that of *recovery*. According to this postulate, if a removed sentence is reinserted into the contracted belief set, then the original belief set is recovered, or, more precisely, can be recovered by logical closure:

$$A \subseteq \text{Cn}((A - \alpha) \cup \{\alpha\})$$

The plausibility of recovery—especially in situations where $A \neq \text{Cn}(A)$ —has repeatedly been questioned (see e.g. Makinson (1987), Fuhrmann (1991), Hansson (1991), Levi (1991)). Accordingly, in AGM (1985) recovery is stated for theories only.

When generalising contraction functions to apply to open belief sets one cannot simply discard recovery without offering a condition to take its place. For without recovery the remaining AGM postulates do not suffice to achieve minimality of belief change. Whereas the inclusion condition precludes the addition of new sentences and the vacuity condition rules out retractions when none are necessary, recovery is the only one among the AGM postulates that prevents unmotivated retractions in the general case. It does so by ensuring that incisions into a theory are kept so small that contractions can be undone by simply adding the removed sentence. Indeed, without recovery the other five postulates are compatible with an operation of theory “withdrawal” (Makinson 1987) such that if $\alpha \in T$, then $T - \alpha = \text{Cn}(\emptyset)$; see Hansson (1991). Thus, we need some postulate that imposes informational economy.

A condition that, like recovery, requires contraction to treasure information is the postulate of *relevance*. Its singleton version was introduced in

Hansson (1991). Fuhrmann and Hansson (1994) extend singleton relevance to the multiple case. Informally, relevance gives expression to the idea that in contracting one should not remove items without reason: whatever is being removed from a belief set in the course of a contraction does in some way contribute to entailing the sentence to be retracted. Relevance for singleton contraction is defined as follows:

- If $\alpha \in A \setminus (A - \beta)$, then there is some set A' such that
- (a) $A - \beta \subseteq A' \subseteq A$,
 - (b) $A' \not\vdash \beta$ and
 - (c) $A', \alpha \vdash \beta$

When A is a theory, then relevance and recovery are equivalent in the presence of the other basic AGM postulates; see Fuhrmann and Hansson (1994), Observation 2. For open A , however, relevance and recovery come apart in both directions of entailment.

The singleton version of relevance generalises to apply to package and choice contractions as follows.

- If $\alpha \in A \setminus (A \multimap B)$, then there is some A' such that
- (a) $A \multimap B \subseteq A' \subseteq A$,
 - (b) $A' \not\vdash B$ and
 - (c) $A', \alpha \vdash B$ P-relevance

- If $\alpha \in A \setminus (A \multimap B)$, then there is some A' such that
- (a) $A \multimap B \subseteq A' \subseteq A$,
 - (b) $A' \not\vdash B$ and
 - (c) $A', \alpha \vdash B$ C-relevance

Relevance, though at first sight slightly unwieldy, is on reflection a much more immediate formulation of the maxim of minimal mutilation than recovery. It is also a rather powerful principle which allows for a very compact presentation of the theory of general contractions. For one, the condition of vacuity is derivable from relevance.

LEMMA 3.1

1. *P-relevance implies P-vacuity, and*
2. *C-relevance implies C-vacuity.*

PROOF. It will suffice to give the argument for 1; the argument for 2 is similar. Assume that $A \not\vdash B$ and—for reductio—that there is some sentence α in A which is not in $A \multimap B$. Then we may apply P-relevance to infer the existence of some set $A' \subseteq A$ such that $A', \alpha \vdash B$. But then $A' \cup \{\alpha\} \subseteq A$ whence $A \vdash B$ by weakening—contradiction. ■

Though it is no doubt good to know that general contractions observe the very plausible vacuity condition, it turns out that—in contrast to the corresponding argument in AGM (1985)—we do not need to have this condition at hand when it comes to proving the representation theorems below.

There is a principle closely related to vacuity, which will be useful in a moment. When the set to be retracted from a belief set does not follow (in the sense of either \vdash or \Vdash) from that belief set, the contraction operation will idle; this is the content of the vacuity postulate. There is yet another case in which we would expect contractions to be vacuous. When a set contains theorems of the underlying logic—i.e. consequences of the empty set—, then that set cannot be package-removed from any belief set; and when a set consists of theorems only, no choice can be made so as to choice-remove that set from any belief set—contraction will *fail*. This very plausible constraint on contractions is formulated in the pair

$$\begin{array}{ll} \vdash B \implies A \subseteq A \overset{\forall}{\setminus} B & \text{P-failure} \\ \Vdash B \implies A \subseteq A \overset{\exists}{\setminus} B & \text{C-failure} \end{array}$$

It turns out that both conditions are, like the vacuity conditions, immediate consequences of the respective relevance conditions. As this fact will be particularly useful in proving the representation theorems for general contractions, we record it as a lemma.

LEMMA 3.2

1. *P-relevance implies P-failure, and*
2. *C-relevance implies C-failure.*

PROOF. Again, it will suffice to give the argument for package contraction as the argument for choice contraction is completely similar. Assume P-relevance and

$$(1) \vdash B \quad \text{while} \quad (2) A \not\subseteq A \overset{\forall}{\setminus} B$$

From (1) we obtain by weakening $C \vdash B$ for any C . It follows from (2) that there is some sentence α such that $\alpha \in A \setminus A \overset{\forall}{\setminus} B$. Hence, by relevance, there must be some set A' between $A \overset{\forall}{\setminus} B$ and A such that $A' \not\vdash B$ —contradiction. ■

Vacuity and failure combine to specify exactly conditions under which a contraction operation degenerates to the identity operation:

OBSERVATION 3.3

1. *P-Relevance, P-inclusion and P-success entail*

$$A = A \overset{\forall}{\setminus} B \iff \vdash B \text{ or } A \not\vdash B \quad \text{P-identity}$$

2. *C-Relevance, C-inclusion and C-success entail*

$$A = A \equiv B \iff \Vdash B \text{ or } A \not\Vdash B \quad \text{C-identity}$$

PROOF. It will be enough to give the argument for 1 only.

(\Leftarrow) Immediate from P-vacuity, P-failure (see the last two propositions) and P-inclusion.

(\Rightarrow) Assume for reductio

$$A = A \not\equiv B \quad (1) \quad \not\vdash B \quad (2) \quad A \vdash B \quad (3)$$

It follows from (2) by P-success that $A \not\equiv B \not\vdash B$. So by (1), $A \not\vdash B$ contradicting (3). \blacksquare

It remains to find a proper generalisation of Gärdenfors' *congruence* condition (C5). This postulate ensures that the contraction of a theory by two logically equivalent sentences yields the same result:

$$\alpha \dashv\vdash \beta \implies A - \alpha = A - \beta.$$

At first sight it may be tempting to generalise this postulate to the requirement that a contraction by logically equivalent *sets* yields the same result, i.e.:

$$\text{Cn}(B) = \text{Cn}(C) \implies A - B = A - C \quad (*)$$

The principle will be shown correct for choice contractions, i.e.

$$\text{Cn}(B) = \text{Cn}(C) \implies A \equiv B = A \equiv C \quad \text{C-congruence}$$

The condition (*), however, is not the right kind of congruence condition for package contractions. For $\text{Cn}(p \wedge q) = \text{Cn}(p \wedge q, p)$. Yet, if we put $A = \text{Cn}(p)$ we would have $A \not\equiv \{p \wedge q\} = A$ by vacuity whilst $A \not\equiv \{p \wedge q, p\} \neq A$ by success.

The counterexample to (*) indicates that it is not logical equivalence *simpliciter* that should count in the formulation of the congruence condition but logical equivalence modulo the belief set to be contracted. To capture this idea we define two sets B and C to be *equivalent modulo* A just in case every subset of A that entails (part of) B entails (part of) C and vice versa:

$$B \equiv_A C \iff \forall X \subseteq A : X \vdash B \Leftrightarrow X \vdash C$$

With this definition at hand we can formulate the sought condition as follows:

$$B \equiv_A C \implies A \not\equiv B = A \not\equiv C \quad \text{P-congruence}$$

The principle is plausible for the intended general notion of contraction. Like the other postulates, however, it will only find its full justification in the representation theorem proved below.

The congruence condition was first formulated in Hansson (1992a) under the name of ‘uniformity’; it reappeared under that name in Fuhrmann and Hansson (1994). Note that package congruence makes (*) hold when it should hold. For since $\alpha \dashv\vdash \beta$ implies $\alpha \equiv_A \beta$, for arbitrary A , package congruence entails the AGM congruence condition (C5). Even more obviously does choice congruence reduce to (C5), if contraction by singletons is at issue.

It is now time to take stock. For convenience and future reference the general contraction postulates are collected in the table below.

| Package Contractions | |
|--|------------------|
| $B \equiv_A C \implies A \forall B = A \forall C$ | congruence (PC1) |
| $A \forall B \vdash B \implies \vdash B$ | success (PC2) |
| $A \forall B \subseteq A$ | inclusion (PC3) |
| $\alpha \in A \setminus A \forall B \implies \exists A': A \forall B \subseteq A' \subseteq A \ \& \ A' \not\vdash B \ \& \ A', \alpha \vdash B$ | relevance (PC4) |
| Choice Contractions | |
| $\text{Cn}(B) = \text{Cn}(C) \implies A \exists B = A \exists C$ | congruence (CC1) |
| $A \exists B \vdash B \implies \vdash B$ | success (CC2) |
| $A \exists B \subseteq A$ | inclusion (CC3) |
| $\alpha \in A \setminus A \exists B \implies \exists A': A \exists B \subseteq A' \subseteq A \ \& \ A' \not\vdash B \ \& \ A', \alpha \vdash B$ | relevance (CC4) |

Table: Postulates for General Contractions

The two sets of postulates are more compact than the AGM set of conditions on theory contraction by singletons. Only success and inclusion have remained “roughly unchanged”. Recovery has been replaced by relevance, whereupon vacuity became redundant.

The closure condition (C5) has been waived without replacement. This is because a general closure requirement to the effect that the result of

a contraction ought to be a theory would obviously offend the declared generality of the present approach to contraction. It can be shown, however, that our generalised contraction functions deliver theories when they should, viz when they are given a theory as input.

OBSERVATION 3.4 *For every theory A ,*

1. *P -inclusion and P -relevance imply that $A \overset{\forall}{\dashv} B$ is a theory, for any set B ;*
2. *C -inclusion and C -relevance imply that $A \overset{\exists}{\dashv} B$ is a theory, for any set B .*

PROOF. We only prove 1 as the argument for 2 is similar. Suppose for *reductio* that

$$A = \text{Cn}(A) \quad (1) \quad \text{and} \quad A \overset{\forall}{\dashv} B \vdash \alpha \quad (2) \quad \text{while} \quad \alpha \notin A \overset{\forall}{\dashv} B \quad (3)$$

From (2) we infer by inclusion (PC3) and weakening for \vdash that $A \vdash \alpha$ whence $\alpha \in A$ by (1). So, by (3), $\alpha \in A \setminus A \overset{\forall}{\dashv} B$. We may now apply relevance (PC4) to infer the existence of some set A' such that

$$A \overset{\forall}{\dashv} B \subseteq A' \subseteq A \text{ and } A' \not\vdash B \text{ and } A', \alpha \vdash B \quad (4)$$

From (2) and the third conjunct of (4) we obtain by cut that $A', A \overset{\forall}{\dashv} B \vdash B$. Then it follows from the first conjunct by weakening that $A' \vdash B$, contrary to the second conjunct of (4). ■

Let us regress, for a moment, to the more restrictive format of contraction treated in the AGM theory by considering the fragment of the general theory that pertains to the retraction of singleton sets from closed belief sets.

We observe, first, that the choice and package postulates coincide for singleton sets on the right-hand-side of the contraction operation. So for the fragment under consideration it is of no importance whether we operate with the PC- or with the CC-postulates. This coincidence of package and choice contraction for singletons renders the ambiguity in the notation $A - \alpha$ harmless: whether $A - \alpha$ stands for $A \overset{\forall}{\dashv} \alpha$ or for $A \overset{\exists}{\dashv} \alpha$ makes no difference; in each case the same set is being referred to.

Let CC stand for the set of CC-postulates and let C stand for the set of basic AGM postulates governing contraction function. Then we observe, second:

OBSERVATION 3.5 *CC entails C.*

PROOF. The postulates C2, C3 and C5 follow immediately from the corresponding CC-postulates. For C1 use Observation 3.4.2, for C5 use

Lemma 3.1.2. It remains to show that recovery (C6) can be derived from CC. Indeed, (C6) can be derived from relevance (CC4) alone:

Assume $A = \text{Cn}(A)$ and suppose for reductio that $\beta \in A$ while $\beta \notin \text{Cn}((A - \alpha) \cup \{\alpha\})$. Then

$$\alpha \rightarrow \beta \in A \quad \text{and} \quad \alpha \rightarrow \beta \notin A - \alpha$$

We may now apply (CC4) to infer the existence of some set A' such that

$$A' \not\vdash \alpha \quad (1) \quad \text{and} \quad A', \alpha \rightarrow \beta \vdash \alpha \quad (2)$$

From (2) it follows that $A' \vdash (\alpha \rightarrow \beta) \rightarrow \alpha$ whence, by $(\alpha \rightarrow \beta) \rightarrow \alpha \vdash \alpha$, $A' \vdash \alpha$ —contrary to (1). ■

Let PC stand for the set of postulates for package contraction. Given the coincidence of package with choice contractions for singletons, our observation yields an immediate

COROLLARY. *PC entails C.*

What about the converse? Do the C-postulates entail the CC- (or PC-) postulates when the latter are restricted to theory contraction by single sentences? The answer is Yes. (Let the label CC^r stand for the restricted set of CC-postulates.)

OBSERVATION 3.6 *C entails CC^r .*

PROOF. The postulates (CC1–3) follow immediately from their corresponding C-postulates (C5), (C2) and (C3) respectively. To prove relevance (CC4) we shall use closure (C1), vacuity (C4) and recovery (C6).

Assume, for an arbitrary closed belief set A ,

$$\beta \in A \quad \text{and} \quad \beta \notin A - \alpha \quad (*)$$

We need to find a set A' such that (a) $A - \alpha \subseteq A' \subseteq A$, (b) $A' \not\vdash \alpha$, and (c) $A', \beta \vdash \alpha$. Let

$$A' = A - \alpha \cup \{\beta \rightarrow \alpha\}$$

We may assume that $\alpha \in A$. For otherwise vacuity (C4) would imply $A \subseteq A - \alpha$, contrary to (*).

Clearly, $A - \alpha \subseteq A'$ by the definition of A' . Moreover, $A - \alpha \subseteq A$ by inclusion (C3) and, since $\alpha \in A$, $\beta \rightarrow \alpha \in A$. So $A' \subseteq A$, concluding the verification of (a).

To show (b) assume for reductio that $A - \alpha, \beta \rightarrow \alpha \vdash \alpha$, i.e.

$$A - \alpha \vdash (\beta \rightarrow \alpha) \rightarrow \alpha \quad (1)$$

Since $\beta \in A$ it follows by recovery (C6) that

$$A - \alpha \vdash \alpha \rightarrow \beta \quad (2)$$

From (1) and (2) we obtain by

$$(\beta \rightarrow \alpha) \rightarrow \alpha, \alpha \rightarrow \beta \vdash \beta$$

that

$$A - \alpha \vdash \beta$$

So, by closure (C1), $\beta \in A - \alpha$ —in contradiction to (*).

Claim (c) is an immediate consequence of the definition of A' . ■

COROLLARY. *C entails PC^r , the set of PC-postulates restricted to contraction of theories by single sentences.*

Observe that in verifying the last two observations we had to use resources that go well beyond elementary structural properties of consequence operations and relations. This is in stark contrast to all other propositions proved in this chapter. The arguments for Observations 3.5 and 3.6 did not only presuppose the presence of an implication connective in the language but also non-trivial properties of implication, such as the deduction equivalence

$$A, \alpha \vdash \beta \iff A \vdash \alpha \rightarrow \beta$$

which, together with weakening for \vdash yields

$$A \vdash \beta \implies A \vdash \alpha \rightarrow \beta$$

and the two principles

$$(\alpha \rightarrow \beta) \rightarrow \alpha \vdash \alpha \quad \text{and} \quad (\beta \rightarrow \alpha) \rightarrow \alpha, \alpha \rightarrow \beta \vdash \beta$$

One of the hallmarks of recovery (C6) is that such substantial logical assumptions are needed for verifying it. This sets recovery apart from the other conditions on contraction function and is part of the reason why recovery has frequently been viewed with a good measure of suspicion.

3.7 Tools for constructing general contraction functions

So far we have rested content with an indirect characterisation of contraction operations on belief sets: we have discussed a number of conditions on such functions, mainly by way of generalising the AGM set of basic postulates. We shall now attempt a more direct characterisation: given two sets we shall give recipes for how to construct the contraction of the one set by the other. The postulates will be used as integrity constraints on such constructions in two ways. First, the constructed functions will be required to satisfy the postulates; and, second, every operation satisfying the postulates will have to be definable by means of the proposed construction. Accordingly, the main technical results in the rest of this chapter are representation theorems for general contractions.

We begin by introducing the formal tools that will be used in the next sections to construct models of package and choice contraction.

Corresponding to the distinction between package and choice contraction, there are two ways of plausibly generalising the remainder operation $(\)\perp(\)$ introduced in Section 2.4 (Definition 2.10) to operate on two arbitrary sets instead of one closed set and a singleton set. One of these generalisations has already appeared but little investigated in the literature on belief change. In one of the first papers on the subject Alchourrón and Makinson (1981) defined the set of remainders of a set A after removing a set B as the collection of all maximal subsets of A that do not overlap with B . (They also observed that if A is a closed under a consequence operation, then every remainder of A will be thus closed too.) Extending the notation introduced in Section 2.4 we let $A\perp B$ stand for the set of remainders of A after removing B . (For $A\perp B$ read: “A perp B”.)

DEFINITION 3.7 (Package remainders)

$X \in A\perp B$ if and only if

- (a) $X \subseteq A$,
- (b) $X \not\subseteq B$, and
- (c) $\forall Y : X \subset Y \subseteq A \implies Y \subseteq B$.

For choice contractions we need to consider a different family of remainders. Now we should not maximise on the property of excluding all of the set to be removed but on excluding only some of it. That is, we should consider the set of maximal subsets of A that do not *contain* the set B . (For $A\angle B$ read: “A angle B”.)

DEFINITION 3.8 (Choice remainders)

$X \in A\angle B$ if and only if

- (a) $X \subseteq A$,

- (b) $X \not\vdash B$, and
- (c) $\forall Y : X \subset Y \subseteq A \implies Y \Vdash B$.

LEMMA 3.9 *For all sets A and B of formulae:*

1. $A \perp B = \emptyset$ if and only if $\vdash B$;
2. $A \angle B = \emptyset$ if and only if $\Vdash B$.

PROOF. As to 1, suppose $\not\vdash B$. Then there is some set $X' \subseteq A$ such that $X' \not\vdash B$ (we may let $X' = \emptyset$). Then X' can be extended to a set $X \subseteq A$ such that $X \not\vdash B$ maximally, whence $A \perp B$ is nonempty. Conversely, if $\vdash B$, then no set X can satisfy condition (b) of Definition 3.7. Hence, $A \perp B$ must be empty. The second assertion is proved similarly with \Vdash replacing \vdash in the argument just given. ■

The remainder operations leave us with a family of sets to choose from. Thus, we need some means of representing such choice. For that purpose a selection function is defined which picks out “preferred” remainders.

DEFINITION 3.10 A *selection function* (for a set A) is any function

$$s_A : \wp(\wp(A)) \longrightarrow \wp(\wp(A))$$

such that $\emptyset \subset s_A(\mathcal{X}) \subseteq \mathcal{X}$ if \mathcal{X} is nonempty, and $s_A(\mathcal{X}) = \{A\}$ otherwise. (In the sequel we omit subscripts to s wherever convenient.)

Thus, in the principal case when the collection \mathcal{X} to choose from is nonempty, the function will select some subset of \mathcal{X} . When \mathcal{X} is empty, which will happen exactly in the case governed by the failure condition, s_A will fall back to selecting A .

3.8 Constructing contractions

Perp-based partial meet contractions. Only a slight alteration of the Definition 2.12 of a partial meet contraction over theories is needed to cover the case of package contraction:

DEFINITION 3.11 An operation $\curlywedge: \wp(Fml) \times \wp(Fml) \longrightarrow \wp(Fml)$ is a *\perp -based partial meet (\perp -pm) contraction* over A if there exists a selection function s_A such that

$$A \curlywedge B = \bigcap s(A \perp B)$$

LEMMA 3.12 *For all sets A, B , and C of formulae,*

$$A \perp B = A \perp C \iff B \equiv_A C$$

PROOF. (\implies) Assume the antecedent and suppose $X \subseteq A$ and $X \vdash B$ while $X \not\vdash C$. Then there exists a set $X' \supseteq X$ such that $X' \in A \perp C$ whence $X' \in A \perp B$ by our assumption. So, since $X' \not\vdash B$ we also have $X \not\vdash B$ —contradiction.

(\impliedby) Suppose

$$B \equiv_A C \quad (*)$$

and $X \in A \perp B$. Then X is a subset of A such that $X \not\vdash B$ and for all $Y \subseteq A$, if $X \subset Y$ then $Y \vdash B$. It follows by (*) that X is a subset of A such that $X \not\vdash C$ and for all $Y \subseteq A$, if $X \subset Y$ then $Y \vdash C$, i.e. $X \in A \perp C$, as required. ■

LEMMA 3.13 *The following are equivalent:*

- (1) *relevance (PC4);*
- (2) *if $\alpha \in A$ and $\forall C(A \multimap B \subseteq C \ \& \ C, \alpha \vdash B \implies C \vdash B)$, then $\alpha \in A \multimap B$;*
- (3) *if $\alpha \in A \setminus A \multimap B$, then $\exists C \in A \perp B$ s.t. $A \multimap B \subseteq C$ and $\alpha \notin C$.*

PROOF. (1) \implies (2): By reshuffling.

(2) \implies (3): Assume $\alpha \in A$ and $\alpha \notin A \multimap B$. It follows from (2) that there is some set X between $A \multimap B$ and A s.t. $X, \alpha \vdash B$ while $X \not\vdash B$. Hence there is an extension C of X such that $C \in A \perp B$. Suppose for *reductio* that $\alpha \in C$. Then $X \cup \{\alpha\} \subseteq C$ whence $C \vdash B$ contrary to $C \in A \perp B$.

(3) \implies (1): Again, assume $\alpha \in A$ and $\alpha \notin A \multimap B$. It follows from (3) that there is some set $C \in A \perp B$ s.t. $A \multimap B \subseteq C$ and $\alpha \notin C$. Since $C \in A \perp B$, $C \not\vdash B$; and since $\alpha \notin C$ and C excludes B maximally, $C, \alpha \vdash B$. ■

The lemma provides with (3) a convenient plug-in to be used in the proof of the next theorem. But it also offers a better understanding of the slightly unwieldy condition of relevance. Relevance is, of course, that one of our conditions on contractions that give formal expression to the maxim of minimal mutilation. (3) is an equivalent way of expressing the condition: If α gets removed from A (in the course of retracting B), then this is so because α cannot always be retained when forming the maximal consistent subsets of A that exclude B .

THEOREM 3.14 (*Hansson 1993.*) *For every belief set A , an operation \multimap is a \perp -pm contraction over A if and only if it satisfies the following conditions for package contractions of A :*

- (PC1) *of congruence,*
- (PC2) *of success,*
- (PC3) *of inclusion, and*
- (PC4) *of relevance.*

PROOF. (\implies): We need to verify the postulates (PC1–4) for \perp -pm contraction over A .

Congruence (PC1) follows immediately from the fact that $B \equiv_A C$ implies $A \perp B = A \perp C$ (Lemma 3.13).

For success (PC2) we need to show that $\bigcap s(A \perp B) \not\vdash B$, if B contains no logical truths. It suffices to note that since, by definition, no remainder entails B , the intersection of a set of remainders cannot entail B —unless $A \perp B = \emptyset$ which is the case if $\vdash B$.

Inclusion (PC3) holds because, by definition, all remainders of A must be contained in A .

To prove relevance (PC4) assume $\alpha \in A$ and $\alpha \notin \bigcap s(A \perp B)$. Then there is some set $C \in s(A \perp B)$ with $\alpha \notin C$. Thus, there is a set C such that $A \not\vdash B \subseteq C \subseteq A$ and $C \not\vdash B$. Moreover, since C is maximally B -excluding, $C, \alpha \vdash B$, as required.

(\impliedby): We assume that $\not\vdash$ satisfies the conditions. We need to show that for each set A there exists a selection function s_A such that for all $B \subseteq Fml$,

$$A \not\vdash B = \bigcap s_A(A \perp B)$$

Let s_A be any mapping from sets of sets of formulae to sets of sets of formulae such that

$$s(A \perp B) = \begin{cases} \{X \in A \perp B : A \not\vdash B \subseteq X\} & \text{if } A \perp B \neq \emptyset \\ \{A\} & \text{otherwise} \end{cases} \quad (\text{s})$$

We need to show that

- (1) s is well-defined, i.e. $A \perp B = A \perp C \implies s(A \perp B) = s(A \perp C)$;
- (2) $s(A \perp B) = \{A\}$ if $A \perp B = \emptyset$, which is immediate from the definition (s);
- (3) $s(A \perp B) \subseteq A \perp B$ if $A \perp B \neq \emptyset$, which is likewise immediate from the definition (s);
- (4) $s(A \perp B) \neq \emptyset$; and
- (5) $A \not\vdash B = \bigcap s(A \perp B)$.

For (1) assume that

$$A \perp B = A \perp C \quad (*)$$

We distinguish two cases.

Case 1: $A \perp B = \emptyset$. Then $A \perp C = \emptyset$ by (*) and so $s(A \perp B) = A = s(A \perp C)$ by (s).

Case 2: $A \perp B \neq \emptyset$. It follows from (s) that

$$X \in s(A \perp B) \implies A \not\vdash B \subseteq X \quad (**)$$

Suppose that $X \in s(A \perp B)$ whence $A \multimap B \subseteq X$ by (**). It follows from (*) by Lemma 3.13 that $B \equiv_A C$ whence $A \multimap B = A \multimap C$ by the congruence condition (PC1). Thus, $A \multimap C \subseteq X$ and so, again by (s), $X \in s(A \perp C)$ as required.

(4) is trivial whenever $A \perp B = \emptyset$. In this case (s) puts $s(A \perp B) = \{A\}$. So consider the principal case, $A \perp B \neq \emptyset$. Then, by Lemma 3.9, $\emptyset \not\vdash B$. So, by success (PC2), $A \multimap B \not\vdash B$. Hence, there is some set C such that $A \multimap B \subseteq C \subseteq A$ and C is maximal in not entailing any of B , i.e. $C \in s(A \perp B)$.

To show (5) we distinguish again two cases.

Case 1: $A \perp B = \emptyset$. Then the definition (s) puts $\bigcap s(A \perp B) = A$ whence the left to right inclusion,

$$A \multimap B \subseteq \bigcap s(A \perp B)$$

follows immediately by inclusion (PC3). For the converse inclusion, note that if $A \perp B = \emptyset$ then $\vdash B$ by Lemma 3.9. So we can apply failure, which follows from relevance (PC4) (Lemma 3.2), to obtain $A \subseteq A \multimap B$.

Case 2: $A \perp B \neq \emptyset$. Since in this case we have

$$X \in s(A \perp B) \iff X \in A \perp B \ \& \ A \multimap B \subseteq X$$

the inclusion $A \multimap B \subseteq \bigcap s(A \perp B)$ follows immediately from (s). To show the converse, assume that

$$\alpha \notin A \multimap B \tag{*}$$

We show the existence of some set $X \in s(A \perp B)$ with $\alpha \notin X$, i.e.: there is some $X \in A \perp B$ with $A \multimap B \subseteq X$ and $\alpha \notin X$. We may assume that

$$\alpha \in A \tag{**}$$

For, if $\alpha \notin A$, then no remainder in $A \perp B$ will contain α and so the required conclusion follows trivially. From (*) and (**) the existence of the required set X follows now essentially by relevance (PC4) via Lemma 3.13. ■

The congruence condition (*) of page 49—the condition that logically equivalent sets determine the same contraction—which is used in the AGM characterisation of singleton partial meet contractions, does not hold for general \perp -pm contraction. But even in the case of singleton contraction there is a sense in which the congruence condition (PC1) rather than (*) should be the preferred postulate. For it just so happens that in the singleton case congruence in the sense of (PC1) can step in for (*). Extension to the general case, however, reveals that (PC1) would have been the more

informative postulate to use all along. This suggestion is based on two observations.

First, if B and C are singleton sets, say $\{\beta\}$ and $\{\gamma\}$, then congruence implies extensionality, $\beta \dashv\vdash \gamma \implies A \forall \{\beta\} = A \forall \{\gamma\}$. Thus, (*) for singletons is correct for general \perp -based p.m. contractions. Second, it is trivial from the definition of \equiv_A that if β and γ are in A and $\beta \equiv_A \gamma$, then $\beta \dashv\vdash \gamma$. It is this result that allows to use the weaker condition (*) rather than congruence in completing step (1) in the above proof when singleton contractions only are at issue; cf. Observation 2.5 of AGM (1985) where the required step (1) is, as elsewhere in the AGM-literature, omitted.

We now come back to the question (posed in Section 3.4) whether package contraction can be reduced to singleton contraction, using the reduction schema

$$A \forall \{\alpha_1, \dots, \alpha_n\} = A - \alpha_1 \cap \dots \cap A - \alpha_n \quad (1)$$

Note that (1) implies

$$B \subseteq C \implies A \forall C \subseteq A \forall B \quad (2)$$

Despite some initial attraction of (2), the partial meet modelling adopted here affords simple counterexamples to (2). Let

$$A = \{p, q, p \rightarrow q\}, \quad B = \{p, q\}, \quad \text{and} \quad C = \{q\}$$

Then $A \perp C = \{\{p\}, \{p \rightarrow q\}\}$ while $A \perp B = \{\{p \rightarrow q\}\}$. Now, we have

$$\bigcap s(\{p\}, \{p \rightarrow q\}) \subseteq \bigcap s(\{p \rightarrow q\})$$

as (2) would require? Only if s selects $\{p \rightarrow q\}$ on the left-hand-side. But, of course, $\{p\}$ may be the only selected remainder in $A \perp C$ in which case (2) would fail.

Angle-based partial meet contractions. Choice contractions can be characterised in a way that corresponds closely to what was obtained for the \perp -operation.

DEFINITION 3.15 An operation $\exists: \wp(Fml) \times \wp(Fml) \longrightarrow \wp(Fml)$ is an *angle-based partial meet (\angle -pm) contraction* over a belief set A , if there exists a selection function s_A such that for each set $B \subseteq Fml$,

$$A \exists B = \bigcap s(A \angle B)$$

LEMMA 3.16 *Let A , B , and C be sets of formulae. Then*

1. if $\text{Cn}(B) = \text{Cn}(C)$, then $A \angle B = A \angle C$, and
2. if $B \subseteq A$ and $C \subseteq A$ and $A \angle B = A \angle C$, then $\text{Cn}(B) = \text{Cn}(C)$.

PROOF. *Ad 1:* We assume that

$$\text{Cn}(B) = \text{Cn}(C) \quad (*)$$

and suppose that $X \in A \angle B$, i.e.

$$X \subseteq A \text{ (a)} \quad \text{and} \quad X \not\vdash B \text{ (b)} \quad \text{and} \quad \forall Y : X \subset Y \subseteq A \implies Y \vdash B \text{ (c)}$$

It suffices to note that (*) allows to replace B by C in (b) and (c). Hence, $X \in A \angle C$ as required.

Ad 2: Assume the antecedent,

$$B, C \subseteq A \text{ (1)} \quad \text{and} \quad A \angle B = A \angle C \text{ (2)}$$

It will suffice to show for any $X \subseteq A$ that $X \vdash B \iff X \vdash C$. For then we have by (1) and the fact $Y \vdash Y$ (for any Y) that $B \vdash C$ and $C \vdash B$, i.e. $\text{Cn}(B) = \text{Cn}(C)$. Thus, suppose for reductio that for some $X \subseteq A$, $X \vdash B$ while $X \not\vdash C$. Then there exists a set X' such that $X \subseteq X'$ and X' excludes (part of) B maximally. Hence, $X' \in A \angle B$ and so, by (2), $X' \in A \angle C$. It follows that $X' \not\vdash C$. But X' contains X whence $X \not\vdash C$ —contradiction. ■

THEOREM 3.17 *For every belief set A , an operation Ξ is an \angle -pm contraction over A if and only if it satisfies the following conditions for choice contractions of A :*

- (CC1) of congruence,
- (CC2) of success,
- (CC3) of inclusion, and
- (CC4) of relevance.

PROOF. The argument parallels that for Theorem 3.14. We leave the proof of sufficiency to the reader.

To prove the necessity direction of the theorem we define for each belief set A a function s_A such that

$$s(A \angle B) = \begin{cases} \{A' \in A \angle B : A \Xi B \subseteq A'\} & \text{if } A \angle A \neq \emptyset \\ \{A\} & \text{otherwise} \end{cases} \quad (\text{s})$$

Again, we need to verify that

- (1) s is well-defined, i.e. $A \angle B = A \angle C \implies s(A \angle B) = s(A \angle C)$;
- (2) $s(A \angle B) = \{A\}$ if $A \angle B = \emptyset$, which is immediate from the definition (s);

- (3) $s(A\angle B) \subseteq A\angle B$ if $A\angle B \neq \emptyset$, which is likewise immediate from the definition;
- (4) $s(A\angle B) \neq \emptyset$; and
- (5) $A \stackrel{\exists}{=} B = \bigcap s(A\angle B)$.

Instead of displaying the complete proof—thereby repeating the pattern of the argument for Theorem 3.14—we just indicate where the four postulates will have to be used.

For (1) we apply Lemma 3.16 and then use congruence (CC1). To prove (4) one uses success (CC2). For (5) we distinguish again two cases. In the case $A\angle B = \emptyset$ the desired conclusion follows by inclusion (CC3). In the principal case, where $A\angle B$ is nonempty, relevance (CC4) is needed. ■

3.9 Reduction strategies revisited

In Section 3.4 some evidence was presented which strongly suggested that both kinds of multiple contractions—package and choice—are independent of singleton contractions. In this section we shall look at reduction strategies more thoroughly. I shall argue that such strategies are not likely to be successful. There remains a remnant of uncertainty, largely because the notion of a reduction strategy is slightly vague. Still, given certain constraints on the range of reduction strategies to be considered, the theory of multiple contractions appears to be not only a natural and in many ways necessary but indeed a proper extension of the theory of singleton contraction.

One version of the reduction thesis is based on the possibility of repeated or iterated contractions. The picture is as follows. There is a space of belief sets. To each belief set A we may associate two families of belief sets: those “accessible” from A by way of a sequence of multiple contractions and those accessible from A by way of a sequence of singleton contractions. The reduction thesis asserts that for any given belief set the two families coincide. In an obvious analogy with the concept of functional equivalence in Boolean sentential logic, this version of the reduction thesis amounts to the thesis that the theory of singleton contraction is *functionally equivalent* to the theory of multiple contractions.

This thesis cannot be assessed in the present framework. For, just like the classical AGM theory of belief change, the present theory yields no general principles concerning *iterated* belief change. But, obviously, such general principles would be needed in any attempt to prove functional equivalence in the above sense.

But even if, on some account of iterated belief change, the thesis turned out true, it would provide for a rather unsatisfactory concept of reduction.

For the question we should like to have an answer to is this: Given a set A , is there a singleton set $\{\alpha\}$ such that retracting A from a belief set has the same effect as retracting $\{\alpha\}$ from that set? (Here we deliberately remain ambiguous as to whether package or choice contraction is at issue.) We expect an answer that specifies the formula α solely in terms of the set A ; what we are after is a recipe for turning any given set into a formula such that the set and the formula are contraction-equivalent. In particular, there should be only one such recipe for turning sets into formulae—not one recipe for one belief set and another recipe for another belief set. Functional equivalence on its own is not very helpful. We should also like to know how to turn information about the composition of some set into information as to which sequence of singleton contractions we should engage in to achieve the same effect as if we retracted that set in one step. This requirement goes well beyond the thesis of functional equivalence and it is doubtful whether it can be satisfied by a plausible theory of iterated contractions.

Let us now restrict our attention to those substantive and helpful versions of the reduction thesis that derive from the following schema:

The Reduction Schema:

Let A be any belief set. For each set $B \subseteq Fml$ there exists a formula fB such that $A - B = A - fB$.

Particular instances or classes of instances of the schema are generated by

- (a) specifying the kind of contraction operation at issue and by
- (b) defining the mapping $f : \wp(Fml) \longrightarrow Fml$.

If nontrivial mappings f can be found for which the thesis is true, then every multiple contraction $A - B$ may be represented by a singleton contraction $A - fB$.

As explained a moment ago, mappings to be avoided are those that depend on the particular contraction at issue. Thus, the equation $A - B = A - fB$ should not be forced to hold by specifying f in terms of A or even $A - B$.

Other conditions are less obvious but still reasonable. For example, the formula fB should be constructed only from material supplied by B , i.e.

$$\text{var}(fB) \subseteq \bigcup \{\text{var } \beta : \beta \in B\}$$

where $\text{var}(fB)$ is the set of variables occurring in the formula fB .

It would be nice, if matters could be kept “simple”, e.g. conforming to the condition

$$fA \text{ is a truth function of a finite subset of } A$$

Without some such substantial constraints the reduction schema is impossible to assess: there are too many—in particular too many gerrymandered—ways of filling in (a) and (b). Let us therefore focus on a reduction thesis which is particularly natural and tempting. This thesis can be factored out into two components:

Finitude:

For each set $B \subseteq Fml$ there exists a finite subset B' of B such that for each belief set A , $A - B = A - B'$.

Sententiality:

For each finite set $B \subseteq Fml$ there exists a formula fB such that $A - B = A - fB$.

The biggest hurdle to reduction is the finitude component: contraction by an infinite set cannot in general be simulated by retracting a finite set.

To see why, note first that there are sets, say A and B , for which choice and meet contractions coincide: $A \forall B = A \exists B$. To take a classical—though finite—example, suppose you decide to open your mind to the possibility that Bizet and Verdi were compatriots. Then it suffices to choice-retract the set {Bizet is French, Verdi is Italian}. But if you cannot make up your mind which one of the two sentences to give up, you will have to retract both. In that case choice and package contractions are indistinguishable.

Consider now a belief set A and an infinite set B such that $A \forall B = A \exists B$. Let all sentences in B be logically independent. Thus, if $\alpha, \beta \in B$, then no logical connection between α and β settles the question whether $\alpha \in A - \beta$ or *vice versa*. Extending this thought, we may observe that for arbitrary subsets B' of B , no logical considerations can settle the question whether $A - B = A - B'$. Thus, A and B may be chosen such that for no — finite or infinite—proper subset B' of B do we have $A - B = A - B'$.

This abstract counterexample to finitude may be fleshed out in many ways. Suppose for example that one decides to retract from A the proposition that Peter hates all prime numbers. Let us also assume that the universal closure of a sentence is in A just in case A contains all of its instances (ω -completeness). So retracting that Peter hates all prime numbers requires the removal of an infinite set of sentences,

$$B = \{ \text{Peter hates the number } n: n \text{ is prime} \}$$

To make the universal sentence fail, it suffices to retract one of the sentences in B . So we have a paradigm case of choice contraction at hand. Moreover, as long as there is no information as to which particular prime number(s)

Peter has ceased to hate, it will be best to withdraw all of B . Thus we have a case where the available information does not suffice to distinguish among the candidates for choice and, hence, forces to discard all candidates; we have $A \stackrel{\exists}{=} B = A \stackrel{\forall}{\neq} B$. Clearly, since all elements in B are logically independent, there is no finite or even infinite proper subset of B whose deletion from A has the same effect as the deletion of all of B . Thus finitude fails.

In contrast to the finitude component, it is rather doubtful whether a general negative result concerning the sententiality component can be obtained: again the range of specifying the sentence-forming function f is too vast. For the time being we rest content with a negative conjecture concerning package contractions: there is no “natural” such function. (We reserve the right to specify necessary conditions of naturalness when faced with unnatural functions.) However, a particularly simple positive observation can be made concerning choice contractions.

OBSERVATION 3.18 *For all finite sets B and any belief set A , $A \stackrel{\exists}{=} B = A - \bigwedge B$.*

PROOF. It suffices to show that $A\mathcal{L}\{\alpha, \beta\} = A\mathcal{L}\{\alpha \wedge \beta\}$, for arbitrary sentences α and β , which follows immediately from the fact that $X \Vdash \alpha, \beta \iff X \Vdash \alpha \wedge \beta$. ■

No corresponding simple relationship holds for package contraction. In particular,

$$A \stackrel{\forall}{\neq} B = A \stackrel{\forall}{\neq} \bigvee B, \text{ for finite } B$$

fails, essentially because $X \Vdash \alpha, \beta$ does not follow from $X \Vdash \alpha \vee \beta$. (Recall the Definition 2.1 of \Vdash and the subsequent discussion.) An exercise in juggling truth-tables reveals, that it does not help to replace \bigvee by some other truth function of elements of B .

To sum up, reduction of either choice or package contraction via finitude is not in general feasible. For finite sets choice contractions may easily be represented as singleton contractions. A dual strategy fails for package contractions.

There are reduction strategies not covered by the above reduction schema. For example, we may allow for representations defined in terms of intersections of sets of singleton contractions, or, as in the first strategy discussed, we may allow for sequences of singleton contractions. Variations and combinations abound: their exploration is left to the adventurous reader.

3.10 Theory contraction through base replacement

Most of the considerations brought forward in the remainder of the present chapter will be neutral—or, rather, harmlessly ambiguous—with respect to the distinction between package and choice contraction. We shall thus help ourselves to the benefit of the simpler notation $A - B$ to stand for both, $A \dot{\simeq} B$ and $A \stackrel{\equiv}{=} B$.

So far we have taken it to be one of the hallmarks of contraction operations that they are mappings to subsets of their arguments. This is the contents of the inclusion postulate

$$A - B \subseteq A \qquad \text{inclusion}$$

In the principal case where A contains (part of) B and B is retractable, $A - B$ will indeed be a proper subset of A .

Now, as long as we confine attention to theories, i.e. belief sets closed under consequence, the inclusion condition is entirely reasonable. If we are to effectively remove some sentence from a theory, then that task can only be accomplished by moving to a weaker theory, and ‘weaker’ in this context must mean ‘smaller’.

But the situation changes when we move to the more general case where the items to be contracted need not be closed under consequence. The principal picture we have in mind is this. Theories are generated from (usually finite) bases by closure under logical consequence. When a theory needs contracting, the contraction is performed by contracting its representing base.

Now, if the inclusion condition is to cover the general notion of contracting a belief set, closed or open, then, in particular, it requires that open belief bases representing theories are always contracted by *erasing* some of their elements, as illustrated in the figure below:




Fig. 3.1: Theory contraction through base contraction

There is a sense, however, in which inclusion can place a too demanding condition on contracting an open belief base as opposed to contracting a

theory. For consider the task of retracting a sentence β from a belief base A . It suffices to adjust A such that β is no longer *entailed* by A . That is to say, for the retraction to be successful, inclusion—the condition that $A - \beta$ is contained in A —is not required; we only need to make sure that the weaker condition

$$\text{Cn}(A - \beta) \subseteq \text{Cn}(A) \quad (*)$$

is satisfied. Of course, if A is a theory, then, as we have seen in the last chapter, $A - \beta$ will also be a theory. Hence, when confining attention to theories, (*) is equivalent to the stronger condition of inclusion. To sum up: it appears that in the most general setting, where retractions from open or closed belief sets are at issue, inclusion is not needed and should be replaced by the weaker condition of *deductive inclusion*:

$$A \Vdash A - B \quad \text{inclusion}^*$$

But not only may inclusion not be needed: sometimes it should not be wanted either. The condition can incur a gratuitous loss of information, thus violating the Maxim of Minimal Mutilation. Consider the following abstract example.

Let your belief base A contain just two sentences, α and β . Suppose you acquire evidence that (1) forces you to give up either α or β but (2) is not specific enough to tell you which one of the two needs to be rejected and which one may be kept. You cannot refuse to retract because one is false for certain. Neither can you make an arbitrary choice for it may be the wrong one. So, in the interest of avoiding error, you need to give up both. That is to say, you are forced to contract A to the empty set. But that seems wrong. For you know that one of α or β is still a valid piece of information. Thus, you should continue accepting the disjunction $\alpha \vee \beta$ after contraction. But then you cannot shrink the base A —you must *replace* it, as illustrated more generally in the following figure:




Fig. 3.2: Theory contraction through base replacement

Counterexamples to inclusion with a different flavour are familiar from “axiom chopping”. Let classical propositional logic be presented by the

following axiom schemas and rule:

$$\alpha \rightarrow (\beta \rightarrow \alpha) \quad (\text{A1})$$

$$(\alpha \rightarrow (\beta \rightarrow \gamma)) \rightarrow ((\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrow \gamma)) \quad (\text{A2})$$

$$\neg\neg\alpha \rightarrow \alpha \quad (\text{A3})$$

$$\frac{\alpha, \alpha \rightarrow \beta}{\beta} \quad (\text{MPP})$$

Few logicians would take the question “What would remain of classical logic if we rejected (A1)?” as an invitation to consider the logic determined by the axioms (A2), (A3) and the rule (MPP). There are presumably many answers to the question (see e.g. Anderson and Belnap (1975)) but at the very least, we should like to leave the Identity schema

$$\alpha \rightarrow \alpha \quad (\text{I})$$

unaffected by the proposed change. As is well-known, the schema (I) is provable from (A1–3) using (MPP); but it is not provable from (A2–3) and (MPP). Whatever logic will be proposed in response to the above question, it will have to be presented by way of a set of axioms that allows a derivation of (I). Hence, that set cannot be a proper subset of the above set of axioms for classical logic but must be a (deductively weaker) *replacement*.

Finally, let us note that the strong inclusion condition is particularly implausible whenever the belief base A is a singleton set, say $\{\alpha\}$. For, in the principal case where $A \vdash B$ (for package contraction) or $A \Vdash B$ (for choice contraction) and B is retractable—that is, not too contaminated with logical truths—, $A - B$ will have to be a proper subset of A . But the only proper subset of a singleton set is the empty set. Thus, in the principal case, we have in general that

$$\{\alpha\} - B = \text{Cn}(\emptyset)$$

which amounts to a trivialisation of contraction operations on singleton belief bases. We shall return in Section 3.12 to the special case of retracting from singleton bases.

3.11 Why general contraction is general enough

In response to the above considerations one may be tempted to adjust the theory of general contraction such that its set of characteristic postulates is replaced by the one below. Let us use the notation $(\dot{-})$ to distinguish the new *contraction* operation from the general contraction operation introduced earlier. (For the sake of definiteness it is at this point advisable to focus on one of the two branches of contractions; we here choose the package branch.)

$$\begin{aligned}
A \equiv A' &\implies A \dot{-} B = A' \dot{-} B && \text{l-congruence (PC1.1)} \\
B \equiv_A B' &\implies A \dot{-} B = A \dot{-} B' && \text{r-congruence (PC1.2)} \\
A \dot{-} B \vdash B &\implies \vdash B && \text{success (PC2)} \\
A \Vdash A \dot{-} B &&& \text{inclusion* (PC3*)} \\
A \vdash \alpha \ \& \ A \dot{-} B \not\vdash \alpha &\implies \exists X: A \dot{-} B \subseteq \text{Cn}(X) \subseteq \text{Cn}(A) \\
&&& \& \ X \not\vdash B \ \& \ X, \alpha \vdash B && \text{relevance* (PC4*)}
\end{aligned}$$

Let us refer to these postulates collectively as the *PC*-conditions* in contrast to the set of PC-conditions introduced earlier. Fortunately, no retreat from the PC- to the PC*-conditions is required. The general theory is general enough. For, contractions $(\dot{-})$ may be sufficiently characterised in terms of contractions $(\dot{\vee})$, as the next observation shows.

OBSERVATION 3.19 *If $A \dot{-} B \equiv \text{Cn}(A) \dot{\vee} B$, then a contraction operation $A \dot{-} ()$ over A satisfies the PC*-conditions just in case a contraction operation $\text{Cn}(A) \dot{\vee} ()$ over $\text{Cn}(A)$ satisfies the PC-conditions.*

PROOF. Elementary. ■

Thus, the bridging principle

$$A \dot{-} B \equiv \text{Cn}(A) \dot{\vee} B \tag{†}$$

is not only an independently plausible constraint on the new operation $\dot{-}$, it also suffices for a complete characterisation of contractions in terms of contractions.

However, (†) does not uniquely characterise contractions; it only does so up to logical equivalence. Is it possible to go further?

To answer this question let us look at (†) from the perspective of partial meet models. It follows from the last observation that

$$A \dot{-} B \equiv \bigcap (s(\text{Cn}(A) \perp B))$$

The right-hand-side of the equivalence may be spelled out as a recipe towards constructing $A \dot{-} B$. First, we consider all maximal subsets of $\text{Cn}(A)$ that do not entail B . Second, we assume that there is a set of in some sense preferred such subsets which gets selected. Third, we take what is common to all such subsets. Since $\text{Cn}(A)$ is closed under consequence, so are all remainders and *a fortiori* every intersection of remainders. Hence, $\bigcap(s(\text{Cn}(A) \perp B))$ is closed under consequence. As a last step it remains to uniquely determine $A \dot{-} B$ by choosing a new belief basis which represents $\bigcap(s(\text{Cn}(A) \perp B))$. Short of this last step we have obtained a characterisation of $A \dot{-} B$ up to equivalence, viz. (†). Should the last step be part of our formal theory of belief change?

This is an instructive example to illustrate a distinction between working and idling parts in formal theories. Suppose we represented the last step in the formal theory by introducing a further selection function, b , which selects a base for $\text{Cn}(A) \dot{-} B$. Then (†) can be “sharpened” to

$$A \dot{-} B = b(\text{Cn}(A) \dot{-} B) \quad (\ddagger)$$

What has been gained by proceeding from (†) to (‡)? So far, very little. We have enriched the expressive power of our theory of contraction (or contraction?) but the new resource has not been used to say very much. Of course, the situation would change, as soon as we formulated and investigated the consequences of plausible and interesting constraints on the choice function b . But are there any such constraints that can be expressed with the resources available to our formal theory? It appears that choice of a representation for a belief set is subject to a perplexing array of contextual factors. No interesting and generally acceptable constraints are forthcoming.

However, the simple fact that the function b is tainted with “pragmatics”—and therefore lacks substantial properties at the appropriate level of abstraction—cannot be the only reason for dispensing with it. For the selection function s introduces similar contextual aspects into the formal theory. Furthermore, no conditions worth mentioning have been imposed on the selection process represented by s ; the function does no more than simply postulating the existence of some selection mechanism. So, why are we inclined to say that s earns a living while b idles?

One difference between the two selection functions is certainly that the need for the one, s , arises fairly early on in the theory, whereas the other, b , leads to a comparatively uninformative extension of an existing theory. The function s is, as it were, an “interior” while b is an “exterior” or “peripheral” operation with respect to our modelling of contraction. The s -function represents an essential ingredient in any theory of belief change and suffices to

get a non-trivial theory off the ground, viz. the theory of general contraction. With the b -function that theory can be extended. But the extension seems to add very little if anything to what is being extended.

We conclude from this little excursion about the value of the b -function that the characterisation up to equivalence (\dagger) of $\dot{-}$ in terms of $\dot{\vee}$ gives sufficient support to the claim that the general theory of contractions is general enough to accommodate apparent violations of inclusion. Contractions of belief sets by replacing their bases can be handled within the general theory of contractions.

3.12 Logical subtraction revisited

The problem. In the last section we have seen how to avert the danger that contracting a singleton belief set degenerates to a trivial limiting case of general contraction. Recall that if $\alpha \vdash B$, then $\{\alpha\} \dot{\vee} B$ will be the empty set, for any contingent α and any completely contingent B .

But, as just argued, we can introduce an operation $\dot{-}$ such that $\{\alpha\} \dot{-} B$ will be equivalent to $\text{Cn}(\alpha) \dot{\vee} B$. The latter is not usually equivalent to the empty set even when $\alpha \dot{\vee} B$ is. Thus, if singleton sets are to be contracted, it is recommended that $\dot{-}$ be used rather than $\dot{\vee}$.

This suggestion sheds new light on an old problem in philosophical logic which has proved remarkably resistant to all previous attempts at a solution. The following passage from Wittgenstein's *Philosophical Investigations* (§621) may serve to illustrate the problem:

“Let us not forget this: when I raise my arm, my arm goes up. And the problem arises: what is left over if I subtract the fact that my arm goes up from the fact that I raise my arm?”

In Jaeger (1973) this passage is taken as a point of departure for searching for a new, hitherto overlooked logical connective: *subtraction*. The problem has been taken up by Hudson (1975) and Humberstone (1981). Jaeger proposes a number of constraints on a satisfactory solution. Hudson offers an appealingly simple solution: subtraction is the converse of material implication. Humberstone rejects Hudson's solution and after examining a number of initially promising approaches concludes that the problem remains open. Since then the subject has disappeared from the logico-philosophical literature.

The passage from Wittgenstein suggests the possibility of a binary operation on propositions which reverses in some sense the effect of conjunction. Just, as the information conveyed by two propositions may be joined by

performing the conjunction-operation, so it should be possible to take away part of the information conveyed by a proposition by performing a subtraction operation.

This train of thoughts finds support in an intriguing picture which seems to pervade all semantical theories. According to the picture, every proposition has a definite *content*. There are a number of theories about the contents of propositions. It is not necessary to adopt any particular one of these theories to make sense of the idea of subtraction. All that is required for the present purpose is to allow for contents to be *structured*: contents has parts. If the contents of a proposition is structured into parts, then these parts may be taken away from the whole of which they are parts. Again, on any account of how the contents of a proposition subdivides, it will presumably emerge that some parts of the contents of some propositions will be such that if they are taken away from that contents, then the remaining contents will again represent a proposition: one that resulted by an operation of subtraction.

In natural language figures of speech corresponding to an operation of subtraction seem ubiquitous. There are at least two salient clusters of examples.

First, there are *negative definitions*:

- (1) A fox is a vixen, except it need not be female.
- (2) Belief is knowledge, though it may neither be justified nor acquired in a proper way nor true.
- (3) A hypothesis is a belief that is not fully accepted.
- (4) A lawlike statement is a law, except it need not be true.
- (5) A gratin is a quiche that is not baked in a shell (Beck et al., 1966, p. 173)

The surface structure of negative definitions may sometimes suggest a conjunctive construal, where the *except*-clause, γ , is conjoined to the first clause, β , of the definiens. But a quick reflection reveals that the suggestion ought to be resisted. A fox is not a vixen that satisfies some further properties; belief is not knowledge plus further constraints; and so on. On a conjunctive construal,

$$\alpha \equiv (\beta \wedge \gamma) \quad (\&)$$

all negative definitions would have a contradictory definiens since the intended effect of conjoining γ is to deny part of what β affirms.

There is no doubt something pathological about examples like (1) and (2)—though very much less so with (3) and (4). It is difficult to conceive of situations where one could attribute to someone a grasp of the concept of a

vixen without thereby attributing to her a grasp of the concepts of *female* and *fox*. Similarly, there is no account of knowledge in the philosophical literature that does not presuppose some notion of belief. It is universally agreed that the notion of belief is relatively less problematic than the notion of knowledge. Thus, there would be little point in trying to define—by subtraction or otherwise—belief in terms of knowledge. Some concepts, it appears, have components so salient that the situation cannot realistically arise in which the concept is understood without a grasp of the salient component.

These considerations explain why subtractions, as occur in (1) and (2), are usually of little value and, indeed, strange. We do understand the example sentences but we find it difficult to conceive of situations in which these sentences carry worthwhile information. Still, it is certainly possible that someone knows what a vixen is and what the distinction between male and female amounts to without knowing what a fox is. Such a person may very well benefit from the negative definition (1), as students of philosophy of science ought to benefit from (4) or readers of menus may find (5) helpful.

A second cluster of examples may be collected under the title *universal quantifications admitting exceptions*:

- (6) Everyone wears a hat, except possibly Bob.
- (7) All birds fly, except possibly penguins.
- (8) Every leader will attend the summit, except possibly Yeltsin.
- (9) Everyone agreed; only John has not yet responded.

Such locutions come natural. But they are difficult to make sense of on any conjunctive construal (&). Once it is affirmed that everyone wears a hat, it is affirmed that Bob wears a hat. Someone who then conjoins a denial that Bob wears a hat sounds significantly more confused than we would expect the average utterer of (6) to be. Universal quantification qualified by exceptions is a common and usually well understood figure of speech.

Note that the truth of $\alpha - \beta$ (in this section $-$ will denote the subtraction connective) does not require the falsehood of β “after” subtraction. The truth-requirements for β are to be taken away from α without at the same time adding those for $\neg\beta$. This is the reason why all except-clauses in the examples above are qualified by “possibly” or similar modals. For example, (6) does not imply that Bob goes hatless. Similarly, (7) says no more than that although penguins are birds, one may not infer in general that penguins fly. Sentence (7) does not license the inference from *Tweety is a penguin* to *Tweety does not fly*. To be sure, there is no reason why one should not also consider a subtraction operation with an “additive component” such that the subtraction implies the negation of what has been subtracted. However,

here we decide to focus on the more basic notion of pure subtraction. “More basic”, because a subtraction with an additive component will presumably be definable as $(\alpha - \beta) \wedge \neg\beta$.

Subtraction as sentential contraction. Guided by the part-whole structure of semantic content Jaeger (1973) proposes a number of conditions which a subtraction connective ought to satisfy.

- (J1) Subtraction is functional (up to equivalence). So if α and α' are equivalent and β and β' are equivalent, then $\alpha - \beta$ should be equivalent to $\alpha' - \beta'$.
- (J2) What is subtracted (β) cannot be part of what is left over ($\alpha - \beta$). So $\alpha - \beta$ should not imply β .
- (J3) What is left ($\alpha - \beta$) over will be a part of the original whole (α). So α should imply $\alpha - \beta$.
- (J4) The whole (α) is equal to the sum of its parts ($\alpha - \beta$ and β respectively). So α should be equivalent to the conjunction of $\alpha - \beta$ and β .
- (J5) What is left over ($\alpha - \beta$) cannot be part of what is subtracted (β). So β should not imply $\alpha - \beta$.

(For reasons to become apparent in a moment, I have changed the order in which Jaeger (1973, p. 320) presents these conditions.)

The reader who has followed us so far will not have failed to recognise old acquaintances in a barely new garb. For example, (J1) corresponds to the congruence condition for contractions; (J2) is just one step away from success; (J3) closely resembles weak inclusion; and (J4) combines (J3) with recovery. Only the condition (J5) has no immediate analogue in the theory of contractions.

To bring out the analogy more explicitly let us recap the postulates for the contraction operation $\dot{-}$ (on p. 67) specialised to the case of retracting singleton sets from singleton bases (omitting now set-brackets around singleton elements and writing again $-$ for $\dot{-}$):

$$\alpha \equiv \alpha' \ \& \ \beta \equiv \beta' \implies \alpha - \beta = \alpha' - \beta' \tag{S1}$$

$$\not\vdash \beta \implies \alpha - \beta \not\vdash \beta \tag{S2}$$

$$\alpha \vdash \alpha - \beta \tag{S3}$$

$$\alpha - \beta \vdash \beta \rightarrow \alpha \tag{S4}$$

$$\alpha \not\vdash \beta \implies \alpha - \beta \vdash \alpha \tag{S5}$$

Relevance has been replaced by recovery (S4) and vacuity (S5). This move is admissible by Observations 3.5 and 3.6.

After having brought out the structural similarities between the J- and the S-conditions, let us now turn to the differences.

First, (J2) suffers from insufficient attention to a limiting case. If β is a logical theorem, then anything will imply β and, hence, in particular we shall have $\alpha - \beta \implies \beta$. Thus, condition (J2) needs the restriction expressed in (S2).

Second, there is the condition (J5) which has no analogue among the S-postulates. This is as it should be. For there clearly are instances of the schema $\beta \vdash \alpha - \beta$: let β be \perp or let β be α . The latter example shows that (J5) is not compatible with (J3) and, hence, should be dropped.

Third, there appears to be no analogue to the vacuity postulate (S5) among the J-postulates. This is so because Jaeger presents his conditions within the scope of a restriction to the principal case that β is “contained” in α , i.e. α implies β . Jaeger does not consider the case of α not implying β . From the point of view of the part-whole picture, the condition (S5) handles this case in a plausible way.

The S-conditions look like a good candidates for characterising subtraction. Accordingly, the proposal at hand is to view subtraction as sentential contraction.

Hudson’s proposal. In logic, as elsewhere, attempts to reduce apparent novelty to the well-known count as virtuous. Thus, Jaeger points out that $\beta \rightarrow \alpha$ satisfies all of the J-conditions thereby qualifying as an attractive truth-functional reduction of the subtraction $\alpha - \beta$. But then, there are other truth functions—e.g. $\beta \vee \gamma \rightarrow \alpha$ —which also satisfy the J-conditions. Hudson (1975) argues that among the many truth functions that satisfy the J-conditions, converse material implication occupies a special place ...

“... in that all the others imply it. (The Deduction Theorem assures us that if $\beta \wedge (\alpha - \beta)$ implies α , then $\alpha - \beta$ implies $\beta \rightarrow \alpha$.) And it seems reasonable to stipulate that if there are several different propositions whose conjunction with β is α , then the least (i.e. weakest) of these shall be considered the difference between α and β . This additional stipulation, then, [...] gives us a unique determination of $\alpha - \beta$: it determines that logical subtraction be the converse of material conditionalisation.” (p. 131)

With Hudson’s conclusion “... that the converse of material conditionalisation is the one and only logical notion which deserves the name ‘subtraction’ ” (p. 135) the problem seems to have been laid at rest: no further discussion has appeared in print.

Yet there are reasons for not being satisfied with Hudson's swift solution of the problem. His proposal amounts to strengthening recovery (S4) to an equivalence, i.e. complementing (S4) with

$$(\beta \rightarrow \alpha) \rightarrow (\alpha - \beta) \quad (\text{Hudson})$$

It follows immediately that

$$\neg\beta \rightarrow (\alpha - \beta) \quad (*)$$

Let β stand for *Foxy is female* and let α stand for *Foxy is a vixen*. Why should it follow from *Foxy is male* ($\neg\beta$) that Foxy is a vixen—or, for that matter, anything you like—minus the requirement that Foxy is female? Or, to take an example from the universal quantification with exceptions cluster: From *Bob does not wear a hat* ($\neg\beta$) it certainly does not follow that everyone wears a hat (α) except Bob. The examples that were adduced above and that motivated our interest in the notion of subtraction provide clear counterexamples to Hudson's proposal. Though attractively simple, Hudson's proposal will not do.

A glimpse at possible worlds. Humberstone (1981) too passes a negative verdict on Hudson's proposal on intuitive grounds. He is then left with the problem of providing an alternative understanding of subtraction, a problem which, in the end, he feels forced to leave open.

With the present proposal—subtraction as sentential contraction—a solution is offered in terms of the models for contraction introduced earlier in this chapter. These models do not only provide an understanding of what subtraction can plausibly mean, they also provide a basis for checking the soundness and even completeness of postulates for subtraction.

Humberstone's own attempts at modelling subtraction are in terms of possible worlds. Given a universe W of possible worlds (variables: a, b, c, \dots, x, y, z) what is a necessary and sufficient condition that $\alpha - \beta$ holds at some world a (notation: $a \models \alpha - \beta$)? The question turns out to be considerably more difficult than the simple picture of the part-whole structure of the contents of a sentence would suggest.

In possible worlds semantics, the contents of a sentence α is usually measured by the set of worlds at which α holds or, dually, by the set of worlds that are ruled out by α . Let us here focus on the former (“Californian”) understanding of contents. By a *proposition* we shall now mean any subset of W and we define for each sentence α the proposition associated with (or determined by) α :

$$[\alpha] = \{x \in W : x \models \alpha\}$$

The identification of the contents of a sentence with the proposition determined by that sentence suggests an immediate and simple rendering of the part-whole talk in precise set-theoretic terms: one proposition is part of another proposition just in case the latter set is contained in the former set. According to this proposal, the relations is-part-of and is-contained-in between propositions are converses of each other.

But this is far too simple a theory of how propositions divide into parts. And it is easily seen to be equivalent to Hudson's proposal:

$$\begin{aligned} [\alpha - \beta] &= [\alpha] \setminus [\beta] \\ &= [\beta \rightarrow \alpha] \end{aligned}$$

A more refined theory of the parts of a proposition is needed to yield a sensible account of subtraction. Such a refined theory can be derived from Grove's (1988) possible worlds models of theory change. Although these models are tuned to modelling not contraction but revision operations in the sense of the AGM theory, it is not too difficult to adjust them for the present purpose. In the remainder of this chapter we shall briefly outline how a possible worlds semantics for subtraction can be obtained.

The models

$$(W, \{\Sigma_P: P \subseteq W\}, \models)$$

have as ingredients

- a (nonempty) set W of possible *worlds*;
- for each subset P of W a family Σ_P of subsets of W (called a *system of spheres*) around P ; and
- a *satisfaction* relation \models between worlds and sentences (read $a \models \alpha$ “ α holds at a ” or “ a makes α true”).

Each family Σ_P satisfies four conditions:

- ($\Sigma 1$) *Minimality*: P is the smallest sphere in Σ_P : $P \in \Sigma_P$ and for all $Q \in \Sigma_P$: $P \subseteq Q$;
- ($\Sigma 2$) *Maximality*: W is the largest sphere in Σ_P : $W \in \Sigma_P$; and
- ($\Sigma 3$) *Finite Nesting*: for every $[\alpha] \neq \emptyset$, there are smallest spheres in Σ_P that (non-emptily) intersect $[\alpha]$: if $[\alpha] \neq \emptyset$ then $\exists S \in \Sigma_P$: ($S \cap [\alpha] \neq \emptyset$ and $(\forall S' \in \Sigma_P$: if $S' \cap [\alpha] \neq \emptyset$ then $S' \not\subset S$)

We shall write $\mu_{\Sigma_P}[\alpha]$ (mostly omitting subscripts) to denote the set of smallest spheres in Σ_P that non-emptily intersect the proposition $[\alpha]$. The third condition may then be put thus:

- ($\Sigma 4$) *Finite Nesting*: for every sentence α , if $[\alpha] \neq \emptyset$, then $\mu[\alpha] \neq \emptyset$ in every system of spheres.

The purpose of this condition is to rule out infinite chains of inclusion. It plays the same role as Lewis' limit assumption in his (1973) semantics for counterfactual conditionals. The present systems of spheres differ from those used in Lewis (1973) in two respects.

First, Lewis' systems are centred on a single world—equivalently: the smallest sphere in any system is always a singleton set. This will not do for the present purpose. For, in contrast to Lewis' theory, our points of departure are not single worlds (total states) but sets of worlds (propositions). Our systems of spheres are centered on sets of worlds. The center of such a system of spheres represents the proposition to be subtracted from.

Second, Lewis' systems are, like Grove's and unlike ours, nested: they impose an additional condition of

Connectedness: if $P, Q \in \Sigma_P$ then $P \subseteq Q$ or $Q \subseteq P$

In the presence of Connectedness it follows from $(\Sigma 3)$ that in any system of spheres there will always be a unique smallest sphere that non-emptily intersects $[\alpha]$, for any non-contradictory sentence α . However, to model the basic subtraction postulates above, the condition of connectedness need not be assumed. (A similar construction of spheres that need not be connected is presented in Lindström and Rabinowicz (1991). Their purpose is different from ours, however, as they want to model non-functional belief change.) The figure below presents a pictorial comparison between the three kinds of systems of spheres.

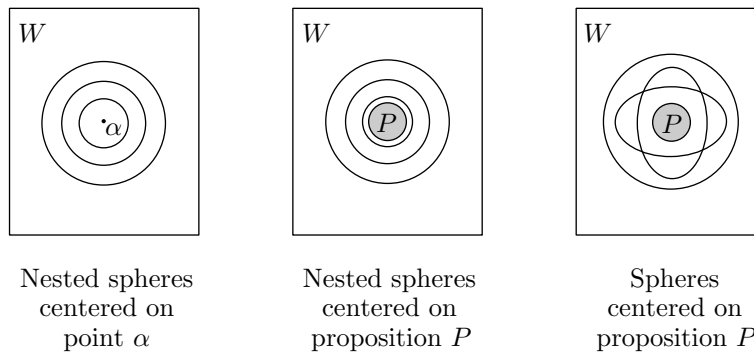


Fig. 3.3

Having defined the frames, let us now turn to the satisfaction relation \models . We assume that \neg and \wedge have their usual Boolean meanings, i.e.

$$[\neg\alpha] = W \setminus [\alpha] \quad \text{and} \quad [\alpha \wedge \beta] = [\alpha] \cap [\beta]$$

We need to find an operation \ominus on propositions such that

$$[\alpha - \beta] = [\alpha] \ominus [\beta]$$

The following validates all S-postulates:

$$P \ominus Q = \begin{cases} P \cup (\bigcup \mu \bar{Q} \cap \bar{Q}) & \text{if } Q \neq W \\ P & \text{otherwise} \end{cases}$$

(where \bar{Q} is the complement of Q in W). Again, the definition may best be understood by way of a pictorial representation:




Fig. 3.4: $[\alpha] \ominus [\beta] = [\alpha] \cup ((S_1 \cup S_2) \cap [\neg\beta])$

The figure exemplifies the principal case in which β is not a logical truth and α implies β . The set of smallest spheres around $[\alpha]$ that successfully reach out to the realm of $\neg\beta$ -worlds consists of just the two spheres S_1 and S_2 . In the picture there are two other spheres—indicated by dotted lines—around $[\alpha]$. But the one dotted sphere does not reach out to the $[\neg\beta]$ area and the other properly includes a sphere, viz. S_1 , which does intersect with $[\neg\beta] = \overline{[\beta]}$. So these two spheres do not qualify, for different reasons, for membership in $\mu\overline{[\beta]}$.

Aided by the figure above it is straightforward to check that all S-postulates come out valid. Moreover, the results of Grove (1988) support the conjecture that the S-postulates *completely* characterise the \ominus -operation in spheres-models.

Revision, Merge and Inference

In this chapter the theory of general contractions will be used to first develop a theory of general revisions and of merges and then to establish a connection between belief change and inference.

Revision and merge are “mixed” change operations involving both contractions and expansions. We shall focus here exclusively on revisions that are generated from general contractions of the package variety. Thus, *contraction* will mean package contraction in this chapter and, accordingly, *revision* will mean package revision, unless, of course, stated otherwise. Merges are insensitive to the distinction between package and choice operations.

Like the theory of general contractions, the theory of general revisions extends to changes of arbitrary sets of sentences by arbitrary sets of sentences. The arguments of a revision operation are thus treated completely symmetrically.

The symmetry of general revisions is, however, of a merely grammatical kind. One fundamental asymmetry remains: The sentences that are to be revised by take priority over the sentences—the belief set—to be revised. This priority finds expression in the success condition for revisions: the set B to be revised by must be fully contained in the result $A*B$ of revising A by B . By contrast, the *merge* operation is a non-prioritised change operation—to borrow a term from Hansson (1991)—which takes both its arguments on equal terms. The sole aim of merge is to blend two sets of sentences to a consistent new set of sentences so as to maximise valuable information. There is no condition here that the information in the one set is to be “more

preserved” than the information in the other set.

We shall proceed by first finding suitable generalisations of the Levi and the Gärdenfors identities to generate revisions from contractions and *vice versa*. Then a theorem will be proved, showing that the theories of revisions and of contractions are interderivable.

Next we turn to merges and display the by now familiar match between an axiomatic characterisation and one in terms of meets of selected remainders. Using the latter we show how the merge operation can be defined either in terms of contractions or in terms of revisions.

In the second part of this chapter we draw out a connection between changes of belief and various styles inference. In particular, so-called *non-monotonic* inference relation, \vdash , failing the monotonicity property

$$\text{if } A \vdash \alpha \text{ then } A' \vdash \alpha, \text{ for all } A \subseteq A'$$

will be of interest here. (In terms of an inference operation, C , monotonicity requires that $C(A) \subseteq C(A')$ if $A \subseteq A'$.) We exhibit two such inference relations relating to revisions and to merges respectively.

4.1 Revision

Let A and B be arbitrary sets of sentences and consider the task of revising A so as to consistently incorporate all of B . In the spirit of the partial meet approach to contractions the following construction looks promising. First, collect all subsets of A that are maximally compatible with B . Then, let $A * B$ result by adding B to what is common to all preferred such subsets of A . More formally, we start by forming a family $A \text{ op } B$ (“ A open for B ”) of sets:

- $X \in A \text{ op } B$ if and only if
- (a) $X \subseteq A$,
 - (b) $X, B \not\vdash \perp$, and
 - (c) $\forall Y : X \subset Y \subseteq A \implies Y, B \vdash \perp$.

Next, we assume that we have some means of selecting from $A \text{ op } B$ a family of preferred sets. As before, selection will be represented by a choice function s . If $s(A \text{ op } B)$ is the set of preferred elements of $A \text{ op } B$, then $\bigcap s(A \text{ op } B)$ consists of just the common elements of all such preferred sets. This is a straightforward adaptation of the partial meet approach to change operations, first presented in Section 2.3, then extended in Sections 3.7–8. Now $\bigcap s(A \text{ op } B)$ is ready for consistently receiving B . It just remains to add B in order to obtain a representation of $A * B$.

Some care needs to be exercised at this point. If, on the one hand, the set A to be revised is a theory, then we should like $A * B$ to be a theory too. If, on the other hand, A is an open belief set, then $A * B$ should not be closed. In general, the operation $*$ should preserve the properties of openness and closedness of belief sets. There should be a “categorical matching” (a term used in Gärdenfors and Rott, 1995) between input states and output states: theories map onto theories and open belief sets map onto open belief sets.

The simplest rendering of *adding* leads to a revision operation that may violate the condition of categorical matching.

$$A * B = \bigcap s(A \text{ op } B) \cup B \quad (1)$$

Here the operation $*$ is defined such that we have a mapping from belief sets to belief sets but not in general one that preserves the property of being closed under consequence. Let us call this operation for the moment being *pre-revision*.

A pre-revision can be thought of as the first step towards a revision operation \star that satisfies the categorical matching condition and is defined thus:

$$A \star B = \begin{cases} \text{Cn}(A * B) & \text{if } A = \text{Cn}(A) \\ A * B & \text{otherwise} \end{cases} \quad (2)$$

Obviously, the non-trivial work in a revision is done by the pre-revision operation. The completion to a full revision is little more than an afterthought that provides for categorical matching in a most straightforward way. For this reason, we shall give prominence in this chapter to the operation of pre-revision, calling it from now on *revision*. We shall briefly return to the \star -operation in Section 4.2.

Some of the properties revisions are expected to have, are immediate from equation (2). For example, the success property $B \subseteq A * B$ follows as immediately as does the inclusion property $A * B \subseteq A \cup B$. Moreover, given that each set in $A \text{ op } B$ is consistent with B , the revision $A * B$ can only be inconsistent, if the set B is itself inconsistent, i.e.: if $A * B \vdash \perp$ then $B \vdash \perp$.

But rather than verifying these and other properties directly, it will be more instructive to see, how the opening operation $() \text{ op } ()$ can be characterised in terms of the (package) remainder operation $() \perp ()$, and how the properties of revision reflect corresponding properties of contractions.

For the purpose of the next lemma recall from Chapter 2 our definition of set-negation:

$$\neg B = \{ \neg \bigwedge B' : B' \subseteq_f B \} \quad (\text{D}\neg)$$

Recall also that this definition delivers the equivalence the definitions of

$(\) \text{op} (\)$ and $(\) \perp (\)$ render $A \text{op} B$ and $A \perp \neg B$ equivalent. Hence we have

$$A * B = \bigcap s(A \text{op} B) \cup B = \bigcap s(A \perp \neg B) \cup B \quad (3)$$

Now, given the representation result for package contractions in terms of partial meets of remainder sets, we have in particular

$$A - \neg B = \bigcap s(A \perp \neg B) \quad (4)$$

whence from (3) and (4) follows the following generalisation of the Levi identity:

$$A * B = (A - \neg B) \cup B \quad (\text{LI})$$

We shall now show how, via the Levi identity, the characteristic conditions for package conditions translate into conditions for revisions. After that we shall prove directly that these conditions are also characteristic, i.e. that for any operation $*$ satisfying the conditions listed in the next theorem a selection function exists such that the partial meet equation (1) holds. We shall thus have shown that (1) applying the Levi identity to an operation of package contraction results in a revision operation (in the sense of the next theorem), and that (2) *every* revision operation can thus be generated from a contraction operation.

LEMMA 4.1 *If $\alpha \in A \setminus A * B$, then the following two conditions are equivalent:*

1. *there is an X such that $(A * B) \cap A \subseteq X \subseteq A$ and $X \not\vdash \neg B$ and $X, \alpha \vdash \neg B$;*
2. *there is an $X \in A \perp \neg B$ such that $(A * B) \cap A \subseteq X$ and $\alpha \notin X$.*

PROOF. $1 \implies 2$: Extend X to a set $X' \in A \perp \neg B$. Then $A * B \cap A \subseteq X \subseteq X'$. Moreover, $\alpha \notin X'$. For, $X', \alpha \vdash \neg B$; hence, if $\alpha \in X'$, then $X' \vdash \neg B$ contrary to $X' \in A \perp \neg B$.

$2 \implies 1$: From $X \in A \perp \neg B$ it follows that $X \subseteq A$ and that $X \not\vdash \neg B$. Moreover, since X is a maximal subset of A not entailing $\neg B$ and since $\alpha \in A \setminus X$, we have $X, \alpha \vdash \neg B$. ■

THEOREM 4.2 *Let $-$ be a contraction operation and let $*$ be the operation generated from $-$ by (LI). Then $*$ satisfies the following conditions.*

- | | | |
|-------------|--|--------------------|
| <i>(R1)</i> | $\neg B \equiv_A \neg C \implies A * B = A * C$ | <i>congruence</i> |
| <i>(R2)</i> | $A * B \vdash \perp \implies \vdash \neg B$ | <i>consistency</i> |
| <i>(R3)</i> | $A * B \subseteq A \cup B$ | <i>inclusion</i> |
| <i>(R4)</i> | $\alpha \in A \setminus A * B \implies \exists X: (A * B) \cap A \subseteq X \subseteq A$ & $X \not\vdash \neg B$ & $X, \alpha \vdash \neg B$ | <i>relevance</i> |

(R5) $B \subseteq A * B$ *success*

PROOF. For *congruence* use the corresponding postulate (PC1) for contractions.

For *consistency* assume $A - \neg B, B \vdash \perp$. Then $A - \neg B \vdash \neg B$ whence we may use success for contractions (PC2) to infer that $\vdash \neg B$.

For *inclusion* we use inclusion (PC3) for contractions, $A - \neg B \subseteq A$, and infer that $(A - \neg B) \cup B \subseteq A \cup B$.

Relevance echoes the corresponding postulate (PC4) for contractions. Assume $\alpha \in A$ and $\alpha \notin A - \neg B$ and $\alpha \notin B$. The consequent of relevance (R4) is just condition 1 of Lemma 4.1. Thus, given the lemma it will suffice to show that condition 2 of that lemma—processed according to (LI)—obtains:

$$\exists X \in A \perp \neg B: ((A - \neg B) \cup B) \cap A \subseteq X \ \& \ \alpha \notin X$$

We may use relevance (PC4), respectively its equivalent according to Lemma 3.13, to infer from our assumption that there is some X such that

$$(1) \ X \in A \perp \neg B, \quad (2) \ A - \neg B \subseteq X, \quad \text{and} \quad (3) \ \alpha \notin X$$

Next we prove that (4) $A \cap B \subseteq X$. Suppose $\beta \in A \cap B$ and—for reductio— $\beta \notin X$. It follows from (1) that any proper extension of X in A will entail $\neg B$; so, in particular, $X, \beta \vdash \neg B$. Thus,

$$X \vdash \neg \beta \vee \neg \bigwedge B_1 \text{ or } X \vdash \neg \beta \vee \neg \bigwedge B_2 \text{ or } \dots \quad (*)$$

for each $B_i \subseteq_f B$ ($i \in \{1, 2, \dots\}$). But for each B_i ,

$$X \vdash \neg \beta \vee \neg \bigwedge B_i \iff X \vdash \neg(\beta \wedge \bigwedge B_i)$$

Since $\beta \in B$, each $\neg(\beta \wedge \bigwedge B_i)$ is itself a member of $\neg B$. So each alternative in (*) entails $X \vdash \neg B$ —contradicting assumption (1). We have thus proved (4).

From (2) and (4) we obtain

$$(A - \neg B) \cup (A \cap B) \subseteq X$$

By inclusion (PC3), $A - \neg B = (A - \neg B) \cap A$ which, by factoring out, gives the required

$$((A - \neg B) \cup B) \cap A \subseteq X$$

■

THEOREM 4.3 *If the operation $() * ()$ satisfies the conditions (R1–5), then there exists a selection function s_A for A such that*

$$A * B = \left(\bigcap s(A \perp \neg B) \right) \cup B$$

PROOF. Let s_A be such that for each set B ,

$$s(A \perp \neg B) = \begin{cases} \{X \in A \perp \neg B : (A * B) \cap A \subseteq X\}, & \text{if } A \perp \neg B \neq \emptyset \\ \{A\}, & \text{otherwise} \end{cases}$$

(i) *The function s is well-defined, i.e. $A \perp \neg B = A \perp \neg C$ entails $s(A \perp \neg B) = s(A \perp \neg C)$.* This is obvious in case $A \perp \neg B$ is empty. For the principal case assume the antecedent. Then by Lemma 3.12, we have $\neg B \equiv \neg C$ whence $A * B = A * C$ by congruence (R1).

(ii) *$s(A \perp \neg B) = \{A\}$ if $A \perp \neg B = \emptyset$ —holds by definition of s .*

(iii) *$s(A \perp \neg B) \subseteq A \perp \neg B$ if $A \perp \neg B \neq \emptyset$ —holds again by definition of s .*

(iv) *$s(A \perp \neg B) \neq \emptyset$.* Obvious (by definition of s) whenever $A \perp \neg B = \emptyset$. So assume that $A \perp \neg B \neq \emptyset$. Then by Lemma 3.9 $\not\vdash \neg B$ whence $A * B \not\vdash \perp$ by consistency (R2). From success (R5) we know that $B \subseteq A * B$. It follows that $A * B \not\vdash \neg B$ and *a fortiori* that $(A * B) \cap A \not\vdash \neg B$. Now, we may extend $(A * B) \cap A$ in A to a set X which among the subsets of A has the property of maximally not entailing $\neg B$. That is to say, X is such that $(A * B) \cap A \subseteq X$ and $X \in A \perp \neg B$.

Assertions (i–iv) amount to saying that s is well-defined. It remains to prove that the equation in the statement of the theorem holds. In case $A \perp \neg B = \emptyset$, we have $\bigcap s(A \perp \neg B) = A$ and we are done immediately. In the principal case $A \perp \neg B \neq \emptyset$ the definition puts

$$X \in s(A \perp \neg B) \iff X \in A \perp \neg B \ \& \ (A * B) \cap A \subseteq X \quad (*)$$

For the left-to-right inclusion of the equation suppose that (1) $\alpha \in A * B$ and (2) $\alpha \notin \bigcap s(A \perp \neg B)$; we show that $\alpha \in B$. It follows from (2) by (*) that there is some set X such that $(A * B) \cap A \subseteq X$ and $\alpha \notin X$ whence $\alpha \notin A * B$ or $\alpha \notin A$. But the first disjunct is impossible in view of (1). So (3) $\alpha \notin A$. From inclusion (R3) we know that $A * B \subseteq A \cup B$. It thus follows from (1) that $\alpha \in A$ or $\alpha \in B$. But the former contradicts (3); hence, $\alpha \in B$.

For the converse inclusion we shall use relevance (R4). Assume $\alpha \notin A * B$. It follows by success (R5), $B \subseteq A * B$, that $\alpha \notin B$. Suppose further that $\alpha \notin A$. Then $\alpha \notin s(A \perp \neg B)$ and we are done. So suppose that $\alpha \in A$. Then it follows by relevance (R4) and Lemma 4.1 that there is some set $X \in A \perp \neg B$ such that $(A * B) \cap A \subseteq X$ (whence by (*) $X \in s(A \perp \neg B)$) and $\alpha \notin X$. ■

COROLLARY 4.4 *An operation $(\) * (\)$ satisfies conditions (R1–5) if and only if $*$ is generated by the Levi identity from a package contraction operation.*

Before reconsidering the question of categorial matching, we record one more property of revisions.

OBSERVATION 4.5 *If $A \not\vdash \neg B$ then $A \cup B = A * B$*

PROOF. The inclusion $A * B \subseteq A \cup B$ is given unconditionally by (R3). It will thus suffice to show

$$A \not\vdash \neg B \implies A \subseteq A * B$$

For, since by success (R5) we have $B \subseteq A * B$, it follows from $A \subseteq A * B$ that $A \cup B \subseteq A * B$. So assume (1) $A \not\vdash \neg B$ and (2) for some α , $\alpha \in A$ while—for contradiction— $\alpha \notin A * B$. It follows from (1) and (2) by relevance (R4) that there is a set $X \subseteq A$ such that $X, \alpha \vdash \neg B$ whence, by weakening, $A, \alpha \vdash \neg B$. Since by (1) $\alpha \in A$, $A \vdash \neg B$, which contradicts (1). ■

As pointed out on page 81, the operation $*$ may violate categorial matching: from $A = \text{Cn}(A)$ it does not follow that $A * B = \text{Cn}(A * B)$. (Recall that both package and choice contractions do satisfy the corresponding condition for contractions.)

Categorial matching can be recovered by closing the output state of a revision process under consequence, if the input state was thus closed. An operation of *matching revision* may be defined thus:

$$A \star B = \begin{cases} \text{Cn}(A * B) & \text{if } A = \text{Cn}(A) \\ A * B & \text{otherwise} \end{cases} \quad (2)$$

The following two observations are easy to verify and show that the operation \star is a straightforward adjustment of the more basic revision operation $*$.

OBSERVATION 4.6 *If $*$ satisfies the conditions (R1–5) on revision operations, then \star satisfies the following conditions.*

- (MR0) $A = \text{Cn}(A) \implies A \star B = \text{Cn}(A \star B)$ *closure*
- (MR1) $\neg B \equiv_A \neg C \implies A \star B = A \star C$ *congruence*
- (MR2) $A \star B \vdash \perp \implies \vdash \neg B$ *consistency*
- (MR3) $A \star B \subseteq \text{Cn}(A \cup B)$ and
 $A \star B \subseteq A \cup B$ if $A \neq \text{Cn}(A)$ *inclusion*
- (MR4) $\alpha \in A \setminus A \star \neg B \implies \exists A': (A \star \neg B) \cap A \subseteq A' \subseteq A$
 $\& A' \not\vdash B$ & $A', \alpha \vdash B$ *relevance*

(MR5) $B \subseteq A \star B$ *success*

OBSERVATION 4.7 *If \star satisfies the conditions (MR0–5) on matching revision operations, then there exists a revision operation $*$ (satisfying (R1–5)) such that*

$$A \star B = \begin{cases} \text{Cn}(A \star B) & \text{if } A = \text{Cn}(A) \\ A \star B & \text{otherwise} \end{cases}$$

Can revisions by finite sets of sentences be reduced to revisions by singletons? Revisions involve a package contraction. In view of the negative reduction result for package contractions one may be excused for rushing to a negative answer to our question—but that would be premature. Consider a simple revision, say

$$A \star \{\alpha, \beta\} = (A - \neg\{\alpha, \beta\}) \cup \{\alpha, \beta\}$$

In general there is no recipe for representing finite package contractions by singleton contractions. But those package contractions that enter (package) revisions are of a special kind: they all and only apply to negated sets. By Lemma 2.4, $\neg\{\alpha, \beta\}$ may be replaced by $\{\neg\alpha \vee \neg\beta\}$. Thus we have

$$A - \neg\{\alpha, \beta\} = A - (\neg\alpha \vee \neg\beta)$$

and therefore

$$\begin{aligned} A \star \{\alpha, \beta\} &= (A - \neg\{\alpha, \beta\}) \cup \{\alpha, \beta\} \\ &= (A - \neg\alpha \vee \neg\beta) \cup \{\alpha, \beta\} \\ &= (A - \neg(\alpha \wedge \beta)) \cup \{\alpha \wedge \beta\} \\ &= A \star \alpha \wedge \beta \end{aligned}$$

Generalising this argument we see that we have proved the following.

LEMMA 4.8 *For every finite set B , $A \star B = A \star \bigwedge B$.*

We close this section with two further observations.

OBSERVATION 4.9 *For all $A, B \subseteq \text{Fml}$:*

1. $(A \star B) \cup (B \star A) = A \cup B$
2. $A \star \top = A - \perp$, *provided A is a theory.*

PROOF. *Ad 1:* By inclusion (R3) we have both $A \star B \subseteq A \cup B$ and $B \star A \subseteq A \cup B$ whence $(A \star B) \cup (B \star A) \subseteq A \cup B$. The converse inclusion follows from success (R5): $B \subseteq A \star B$ and $A \subseteq B \star A$ whence $A \cup B \subseteq (A \star B) \cup (B \star A)$.

Ad 2: We have $A \star \top = (A - \perp) \cup \{\top\}$ by the Levi identity. If A is a theory, so is $A - \perp$ by the closure property of contractions (Observation 3.4.1). Hence, $(A - \perp) \cup \{\top\} = A - \perp$ as required. ■

The last observation suggests that contractions and revisions may not only be used to retract old beliefs, respectively incorporate new pieces of evidence, but also to make consistent (“consolidate”) existing beliefs. According to Observation 4.9.2 there are two ways of achieving consistency: either by retracting *falsum* or by incorporating *verum*. However, in the special case envisaged here, namely when A is a theory, these *consolidation* operations (Hansson, 1991; see also the next section) are of rather little interest. For, suppose A is indeed inconsistent. Then, being closed under classical consequence, $A = Fml$. Hence, $A - \perp$, respectively $A * \top$, will be some maximally consistent set (or the intersection of a number of maximally consistent sets) of formulae, no matter how the inconsistency was injected to the last consistent predecessor of A . If A is not closed, however, inconsistency may be removed in a more discriminating fashion using a merge or consolidation operation. To such an operation and some of its uses we turn now.

4.2 Merge

In a way, merges reflect more directly James’ description of belief change quoted in the introductory chapter. For, when first confronted with a piece of information that resists smooth integration with our present beliefs, we usually try to modify both, old and new information, so as “... to show a minimum of jolt, a maximum of continuity”. Indeed, contrary to the absolute preference for novel information implicit in the idea of revision, James points out that mostly we change beliefs with a preference for old opinions as far as feasible.

“Their influence is absolutely controlling. Loyalty to them is the first principle—in most cases it is the only principle; for by far the most usual way of handling phenomena so novel that they would make for a serious rearrangement of our preconception is to ignore them altogether, or to abuse those who bear witness for them.” (James, 1907, p. 61f.)

Unlike revision, merge opens the possibility that new evidence is partly or even completely ignored if old information is suitably well-entrenched. The merge operation marries old and new information to a consistent whole without giving undue precedence to the one or the other. The result of a merge operation is solely dictated by the informational value carried by its arguments—no bonus for novelty.

Although merges thus capture more directly actual episodes of belief change, they are definable in terms of contractions or revisions. But it is

not possible to take the reverse route, that is, to define contractions and revisions in terms of merges. This asymmetry explains and justifies the precedence we have given to the more abstract operation of contraction. In the light of the results of the previous section we could have given the same precedence to the revision operation. (But see Levi (1980, ch. 3, 1991, ch. 4) for arguments of a different kind marking contractions as more basic than revisions.)

Partial meet merge. A prima facie candidate for the merge of two sets is any deductively maximally consistent subset of their union. For each set C , a set X is a *deductively maximally consistent* (d.m.c) subset of C , if (a) $X \subseteq C$, (b) $X \not\vdash \perp$, and (c) $\forall Y : X \subset Y \subseteq C$ implies $Y \vdash \perp$. It is easy to see that a set X is a d.m.c. subset of C just in case $X \in C \perp \{\perp\}$.

As in the case of modelling the change operations of previous chapters, we shall here pursue a partial meet approach. That is, we assume that, on the one hand, not all d.m.c. subsets of a given set are of equal value to the inquirer; some sets in $C \perp \{\perp\}$ may be more valuable than others. On the other hand, we do not assume that for every set C there is always a unique most valuable set in $C \perp \{\perp\}$. Hence, there must be some way of picking out a distinguished subset of $C \perp \{\perp\}$, for each set C . As before, these assumptions give rise to the following definition of a merge operation.

DEFINITION 4.10 An operation $\circ : \wp(Fml) \times \wp(Fml) \longrightarrow \wp(Fml)$ is a (\perp -based partial meet) *merge operation* over $A, B \subseteq Fml$, if there exists a selection function s_C for $C = A \cup B$ such that

$$A \circ B = \bigcap s_C(C \perp \emptyset)$$

The presupposed definition of a selection function is as in Definition 3.10. Note that for no set C is $C \perp \{\perp\}$ ever empty. The “worst case” obtains when C is maximally inconsistent, i.e. when there is no nonempty consistent subset of C . Thus, the special provision in Definition 3.10 for selecting from the empty set never applies as far as merge operations are concerned.

We now turn to an axiomatic characterisation of merge operations as just defined. (In the proof of the theorem $A \cup B \perp \{\perp\}$ should be read as $(A \cup B) \perp \{\perp\}$, according to the convention that the usual set operations bind more strongly than the remainder operation.)

THEOREM 4.11 *An operation $\circ : \wp(Fml) \times \wp(Fml) \longrightarrow \wp(Fml)$ is a (\perp -pm) merge operation over sets A and B if and only if it satisfies the following conditions:*

- (M1) $A \cup B = A' \cup B' \implies A \circ B = A' \circ B'$ *congruence*
(M2) $A \circ B \not\vdash \perp$ *consistency*
(M3) $A \circ B \subseteq A \cup B$ *inclusion*
(M4) $\alpha \in A \cup B \setminus A \circ B \implies \exists D: A \circ B \subseteq D \subseteq A \cup B$
 $\& D \not\vdash \perp \& D, \alpha \vdash \perp$ *relevance*

PROOF. (\implies) Since $A \cup B = A' \cup B' = C$, $A \cup B \perp \{\perp\} = A' \cup B' \perp \{\perp\} = C \perp \{\perp\}$ whence $s_C(A \cup B \perp \{\perp\}) = s_C(A' \cup B' \perp \{\perp\})$. Thus follows *congruence*. Since all sets in $s(A \cup B \perp \{\perp\})$ are consistent, so is their intersection. Thus follows *consistency*. Since all sets in $s(A \cup B \perp \{\perp\})$ are contained in $A \cup B$, so is their intersection. Thus follows *inclusion*. Finally, to show *relevance* assume that $\alpha \in A \cup B$ while $\alpha \notin A \circ B$. Then there is some $X \in s(A \cup B \perp \{\perp\})$ that does not contain α . Clearly, $A \circ B \subseteq X \subseteq A \cup B$, and since X is maximally consistent, $A \not\vdash \perp$ while $A, \alpha \vdash \perp$.

(\impliedby) Assume that $(\) \circ (\)$ satisfies the conditions (M1–4) for arbitrarily chosen sets A and B . We define a canonical selection function s_C for $C = A \cup B$ to be any mapping $\wp(\wp(Fml)) \longrightarrow \wp(\wp(Fml))$ such that

$$s_C(A \cup B \perp \{\perp\}) = \{X \in A \cup B \perp \{\perp\} : A \circ B \subseteq X\} \quad (s)$$

It needs to be shown that s_C (subscript from now on suppressed) is a well-defined selection function in the sense of Definition 3.10 and that the displayed equation in Definition 4.10 is satisfied.

That s is well-defined follows by *congruence* as follows. Suppose that $X \in s(C \perp \{\perp\})$ and that $C \perp \{\perp\} = C' \perp \{\perp\}$ where $C' = A' \cup B'$. Then $X \in C \perp \{\perp\} = C' \perp \{\perp\}$ and $A \circ B \subseteq X$ by (s) in the left-to-right direction and so $X \in s(C' \perp \{\perp\})$ by applying (s) in the other direction. That $s(C \perp \{\perp\})$ is a subset of $C \perp \{\perp\}$ is immediate from (s). The special case that $C \perp \{\perp\}$ is empty never obtains, as noted above. Hence the stipulation in Definition 3.10 for that case is trivially satisfied. To show that $s(C \perp \{\perp\})$ is nonempty we observe that it follows from *inclusion* and *consistency*,

$$A \circ B \subseteq A \cup B \quad A \circ B \not\vdash \perp$$

that there is a subset of $C = A \cup B$, namely $A \circ B$, which is consistent. Hence, there must exist a set D between C and $A \cup B$ which is maximally consistent. So $D \in C \perp \{\perp\}$ and $A \circ B \subseteq D$, i.e. $D \in s(C \perp \{\perp\})$.

Finally, we show that

$$A \circ B = \bigcap (s(A \cup B \perp \{\perp\}))$$

For the left-to-right inclusion assume that $\alpha \in A \circ B$ and that $X \in s(A \cup B \perp \{\perp\})$. Then $A \circ B \subseteq X$ whence $\alpha \in X$. Since X was chosen arbitrarily, it follows that $\alpha \in \bigcap (s(A \cup B \perp \{\perp\}))$.

For the converse inclusion assume

$$\alpha \notin A \circ B \quad (*)$$

We need to find some $X \in s(A \cup B \perp \{\perp\})$ such that $\alpha \notin X$. We may assume

$$\alpha \in A \cup B \quad (**)$$

for, otherwise no set in $A \cup B \perp \{\perp\}$ —being all subsets of $A \cup B$ —will contain α and so the required conclusion holds immediately. From (*) and (**) we infer by *relevance* the existence of a set D such that

$$A \circ B \subseteq D \subseteq A \cup B \text{ and } D \not\vdash \alpha \text{ while } D, \alpha \vdash \perp$$

whence $D \not\vdash \alpha$. We may extend D to a set $D' \in A \cup B \perp \{\perp\}$. Then

$$A \circ B \subseteq D' \subseteq A \cup B \text{ and } \alpha \notin D'$$

It remains to show that $D' \in A \cup B \perp \{\perp\}$ by verifying that D' is maximally consistent. Suppose that D' can be further extended to a consistent set $D'' \subseteq A \cup B$. Since $D' \in A \cup B \perp \alpha$, we would have $D'' \vdash \alpha$. But $D, \alpha \vdash \perp$ and so $D'' \vdash \perp$ by weakening and cut—contradiction. ■

Consolidation and other definitions of merge. Hansson (1991) has defined on operation () of *partial meet consolidation* on belief bases thus:

$$A! = \bigcap s_A(A \perp \{\perp\})$$

Clearly, we have $A! = A - \perp$. It follows that merge may be defined in terms of consolidation:

$$A \circ B = (A \cup B)!$$

Alternatively, consolidation may be defined in terms of merge:

$$A! = A \circ A$$

Like contraction and revision, merge and consolidation are thus essentially equivalent operations: every change that can be effected by the one operation can also be effected by the other operation. Sometimes it seems more profitable to study changes from the viewpoint of the one rather than the other operation. For example, the unary consolidation operation suggests an analogy with unary modal operators which is drawn out in Hansson (1991). On the other side, the binary merge operation gives rise to a number of structural properties, such as those listed in the next observation.

OBSERVATION 4.12

1. $A \circ B = B \circ A$;
2. $\text{Cn}(A \circ A) \subseteq \text{Cn}(A)$
3. $\text{Cn}(A) \subseteq \text{Cn}(A \circ A)$ if and only if A is consistent;
4. $A \circ (B \cup C) = (A \cup B) \circ C$;
5. $A \circ B = (A \cup B) - \perp$;
6. $A \circ B = ((A \cup B) * \top) \cap (A \cup B)$;
7. $A \circ B \equiv (A \cup B) * \top$;
8. $A \circ A \equiv A * \top$.

The last two equivalences may be strengthened to identities if $A \cup B$, respectively A is a theory.

PROOF. Elementary, using the identities

$$\begin{aligned} A \circ B &= \bigcap s(A \cup B \perp \{\perp\}), \\ C \perp \{\perp\} &= C \perp \{\perp\}, \quad \text{and} \\ C * \top &= \bigcap s(C \perp \{\perp\}) \cup \{\top\} \end{aligned}$$

■

The identities 6 and 7 of Observation 4.12 may serve as definitions of merge in terms of contraction and revision respectively. Both of these definitions offer little surprise; in particular 6 does hardly more than recapitulating the definition of a partial meet merge.

The identity 7 is slightly more involved. In the unlikely case that $A \cup B$ happens to be a theory, we have $A \circ B = (A \cup B) * \top$. In the more interesting case that $A \cup B$ is not a theory, the definition in terms of revision will add *verum* to the result of retracting *falsum* from $A \cup B$. But *verum*—or its defining formula—may not be a member of the open set $A \cup B$ —in which case the addition of *verum* would lead to a proper extension of $A \cup B - \perp$. To cancel the effect of adding *verum* where it does not belong, we intersect with $A \cup B$. The reader may want to convince herself that this is harmless when $A \cup B$ is closed.

Though merge is definable from contractions (and hence also from revisions) the reverse route is blocked: contractions and revisions are not definable in terms of merges. To see why the theory of contractions is essentially more expressive than the theory of merges, consider any of the defining identities of the last observation, say,

$$A \circ B = (A \cup B) - \perp$$

No information about the priority of either A or B can be retrieved from the right-hand-side of the equation. Moreover, merge expressions are clearly

suitable for referring to only a very small subset of contractions, namely those of the form $C - \perp$. It follows that if contractions were definable in terms of merges, all contractions would have to be definable in terms of contracting *falsum*. That is to say, we would expect some *general* reduction scheme of the form

$$A - B = \dots$$

to hold, where the right-hand-side involves only the consolidation-operation $!$ (together with familiar set operations). But simple counterexamples tell against the possibility of such a general schema. For example the contraction $\{p\} - \{p\}$ should be equivalent if not to the empty set, then to some set that contains neither p nor $\neg p$ —no matter how p and $\neg p$ compare in informational or other values on which selection functions may operate. But $\{p, \neg p\}!$, to take the most plausible candidate for a *definiens*, may very well evaluate to p or to $\neg p$. This just reflects the basic fact about consolidation and merges that nothing but inconsistency gets removed for sure—which is not generally a fact about contractions.

4.3 Inference

In the next few paragraphs we shall sketch the notion of nonmonotonic or defeasible inference as the term is used in current research in logic and AI. These remarks serve only to indicate our target notions, not to argue that they are valid abstractions of a cognitive practice that deserves the attention of philosophers and logicians.

For such arguments and a more thorough treatment of the topic the reader should turn elsewhere. Ullman-Margalit (1983) gives a thought-provoking and suggestive “pre-AI” account of what she calls “presumptive inference”. Another rather informal paper but already responding to AI-theories is Bach (1984). Ginsberg (1987/94) is an introduction to varieties of nonmonotonic reasoning from an AI-perspective. (The paper, in its original 1987 version, introduces a collection of seminal contributions to the topic.) Brewka (1991) is a very readable monograph on the subject. A brief orientation is provided in Fuhrmann (199+b). Finally the handbook Gabbay et al. (1994) should be mentioned; it collects a number of excellent expository essays on nonmonotonic and uncertain reasoning. In particular, the handbook contains Makinson’s (1994) impressive study of nonmonotonic inference from a pure logician’s point of view.

In default reasoning one proceeds from explicit premisses to conclusions with the aid of a set of tacit premisses. What distinguishes default reasoning from many other forms of enthymematic reasoning is the caution with which

the tacit premisses, the defaults, are applied. The tacit premisses are such that they can be presumed true under normal circumstances—that is why they “normally” do not need mentioning. But sometimes explicitly supplied premisses may indicate that circumstances are not normal enough to make unrestricted use of all default assumptions. Then some defaults may be bracketed: under the circumstances in question their use would incur an unacceptable risk of error.

Default reasoning, thus characterised, fails to have the following monotonicity property: if some conclusion α follows from a set A of premisses, then α follows from every set that contains A . To see how default reasoning may fail to have this property assume that α follows (by default) from A . Let β be a sentence whose truth would strongly indicate that one or the other presumptions used in jumping from the premisses A to the conclusion α does not hold. Then α does not follow by default from A together with β .

Defaults are sentences that are presumed true unless evidence to the contrary forces one to withdraw one or more of them. Any particular instance of default reasoning draws on two resources: the explicit information given in the particular case (the explicit premisses), and as many of the default (or background) assumptions as are compatible with the explicit information.

The problem faced in modelling default reasoning—referred to as the ‘multiple extensions problem’—is that there is usually more than one way of extending the explicit premisses by as many defaults as are consistent with these premisses. As we have seen in the previous chapters, this problem is not unknown to students of belief change.

Expectation inference

Gärdenfors and Makinson (1994) (“[GM]”) have proposed to solve the multiple extensions problem in default reasoning along the pattern of partial meet constructions of belief changes. Similar proposals had already been made in Gärdenfors (1990) and Lindström (1991). The main achievement of the 1994 paper is its drawing together a large number of *prima facie* diverse modellings of nonmonotonic reasoning and belief change, thereby approximating a uniform theory of both research areas.

The idea is quite simple and the astute reader will not have failed to have guessed the gist of it while perusing the introductory section above. Let D be the set of default assumptions relative to which defeasible inferences are to be performed. Let A be a set of explicitly supplied premisses for the defeasible inference operation C_D . We shall omit the subscript D when

there is no danger of ambiguity. Then the explicit premisses A may be padded with defaults drawn from D as long as consistency is preserved (assuming for the moment that the premise set A is consistent). Each remainder in $D \perp \neg A$ is an (maximally) admissible way in which one may use one's expectations or default assumptions to supplement the premisses A . That is to say, $D \perp \neg A$ collects all subsets of D that may be added to A without engendering inconsistency.

Now we are in a position to define the defeasible inference operation C (depending on defaults D):

$$\beta \in C(A) \iff \beta \in \text{Cn} \left(\bigcap \{X \cup A : X \in s(D \perp \neg A)\} \right) \quad (\text{C1})$$

The function s may now be understood as choosing the best ways of supplementing the premisses with defaults.

Obviously, the right-hand-sides of definition (C1) is strongly reminiscent of the partial meet approach to belief revision. Indeed, it is easily seen that (C1) may equivalently be phrased as a definition of defeasible inference in terms of belief revision:

$$\begin{aligned} \beta \in C(A) &\iff \beta \in \text{Cn}((D - \neg A) \cup A) \\ &\iff \beta \in \text{Cn}(D * A) \end{aligned} \quad (\text{C2})$$

This last equivalence reduces defeasible inference from a set A of premisses to deductive inference from the revision $D * A$ of some fixed set D of defaults. With schema (C2) at hand, the relation between belief revision and defeasible inference can now be studied either at the level of postulates (as in Gärdenfors (1990) or Makinson and Gärdenfors (1991)) or at the level of models (as in Gärdenfors and Makinson (1994)).

There are two deviations from [GM] here. First, the account in [GM] is limited to inference from *finite* premise sets. This is so, because in [GM] revision is defined for singletons only. But a revision by a finite set can be mimicked by a revision by the conjunction of all of its members (Lemma 4.8). Thus, in view of the equation

$$D * A = D * \bigwedge \{\alpha : \alpha \in A\} \quad (A \text{ finite})$$

the account given here coincides with the one in [GM] when finite sets of premisses are at issue.

Second, [GM] reverses the order of taking consequences and intersecting, as in:

$$\beta \in C'(A) \iff \beta \in \bigcap \{\text{Cn}(X \cup A) : X \in s(D \perp \neg A)\} \quad (\text{C'1})$$

In general we have

$$C \leq C'$$

i.e. $\bigcap\{\text{Cn}(X \cap A) : X \in s(D \perp \neg A)\} \subseteq \text{Cn}(\bigcap\{X \cap A : X \in s(D \perp \neg A)\})$.
 But the converse need not hold. However, C and C' coincide when the set of defaults, D , on which both operations depend is closed under consequence—as it is assumed throughout in [GM].

PROOF. ... *that* $C \leq C'$: Without loss of generality we consider the case of only two remainders, X_1 and X_2 . By set theory we have $\text{Cn}((X_1 \cap X_2) \cup A) = \text{Cn}((X_1 \cup A) \cap (X_2 \cup A))$. By properties of Cn the r.h.s. is contained in $\text{Cn}(X_1 \cup A) \cap \text{Cn}(X_2 \cup A)$.

... *that not generally* $C' \leq C$: Let $I = \{1, 2\}$ with X_1 and X_2 nonempty but disjoint (say $\{p_1\}$ and $\{p_2\}$ respectively) and let A be empty. Then $\text{Cn}(X_1) \cap \text{Cn}(X_2) \supset \text{Cn}(\emptyset)$. (For example $p_1 \vee p_2 \in \text{Cn}(X_1) \cap \text{Cn}(X_2)$.)

... *that* $C = C'$ *if* D *is closed*: Note first that if D is closed, then the remainders X_1 and X_2 are also closed. It suffices to show that $\text{Cn}(X_1) \cap \text{Cn}(X_2) = \text{Cn}(X_1 \cap X_2)$. Since X_1 and X_2 are closed, $\text{Cn}(X_1 \cap X_2) = \text{Cn}(\text{Cn } X_1 \cap \text{Cn } X_2)$. Since closure is preserved under intersection, $\text{Cn}(\text{Cn } X_1 \cap \text{Cn } X_2) = \text{Cn}(X_1) \cap \text{Cn}(X_2)$ as required. ■

To sum up: the definition of defeasible inference given here coincides with the one in [GM] for all those cases to which [GM] confines attention: inferences from singletons on the basis of closed sets of defaults. The presentation here extends that of [GM] to inferences from denumerable premise sets on the basis of closed or open sets of defaults.

There are two limiting cases of expectation inference: first, the premisses A may be consistent with the expectations D ; second, A may be inconsistent. In the first case, the set of expectations may be fully relied on in drawing inferences from A . Put more formally: if $\perp \notin \text{Cn}(D \cup A)$, then $D \perp \neg A = \{D\}$ whence $C_D(A) = \text{Cn}(D \cup A)$. This is as it should be.

In the second case, the inconsistent set A will be added to the intersection of the best remainders and the result will be closed under consequence. Hence, $C(A) = \text{Cn}(A) = \text{Fml}$. This is a sorry result for it shows that the defusion of inconsistent information in expectation inference is only half-hearted. Below we shall introduce the concept of merge inference to allow nontrivial defeasible inference from inconsistent premisses.

Poole systems. A particularly simple case of expectation inference is implicit in Poole's (1988) theory of explanation in AI.

So-called Poole systems are both simple and powerful; they have therefore enjoyed considerable attention in the AI-community as a basis for theories of explanation and prediction, causality, diagnosis and abduction; for

a brief survey of these and other areas of potential application see Poole (1994).

Despite its success in the AI-community Poole-systems stand little chance to elicit similar enthusiasm among philosophers. The claims associated with the notion of a Poole-system simply cannot stand scrutiny in the light of a sophisticated philosophical literature on explanation, causality, abduction and related concepts (cf. e.g. Pitt 1988). Poole systems implement a rather crude variant of the deductive-nomological model of explanation. This fact does not detract, however, from the heuristic value of Poole systems, both as toy systems in AI with a limited range of applicability and as points of departure for more elaborate theories. One such a theory, incorporating probabilities, is developed in Poole (1993).

We shall first present the notion of a Poole system using the concepts of Poole (1988) and then point out their simple relation to expectation inference.

A Poole system may be represented as a pair, consisting of a set of *defaults*—Poole uses also the term “possible hypotheses”—and a set of *constraints*. If the latter is empty, a Poole system is called *unconstrained*. The Poole systems considered here will be of the unconstrained kind.

DEFINITION 4.13 For each pair of sets of formulae, A (“facts”) and D (“defaults”), let ...

a *scenario* (of A given D) be a consistent set $X \cup A$ with $X \subseteq D$;

a *maximal scenario* be a scenario no proper extension of which is a scenario;

an *explanation* of a sentence α (from A given D) be a scenario (of A given D) that entails α ;

an *extension* (of A in D) be the closure under consequence of a maximal scenario (of A given D).

Let $e(D, A)$ denote the family of all extensions of A in D , i.e.

$$e(D, A) = \{ \text{Cn}(X \cup A) : X \subseteq D \text{ and } X, A \not\vdash \perp \text{ and} \\ \forall X' \subseteq D : \text{if } X \subset X' \text{ then } X', A \vdash \perp \} \quad (1)$$

Using the definition 3.4 of remainder sets, (1) is equivalent to

$$e(A, D) = \{ \text{Cn}(X \cup A) : X \in D \perp \neg A \} \quad (2)$$

Poole defines a sentence to be *sceptically predictable* from A just in case it is contained in (*explained by*) all extensions of A in D . Plainly, sceptical

prediction, C^{SP} , is a special case of an expectation inference operation as defined in (C'1)—but not as in (C1)!—above:

$$\begin{aligned} C^{\text{SP}}(A) &= \bigcap e(D, A) \\ &= \bigcap \{\text{Cn}(X \cup A) : X \in s(D \perp \neg A)\} \end{aligned}$$

where s now idles in that every remainder gets selected. In the terminology introduced earlier: C^{SP} is an expectation inference operation based on *full meet* contractions.

Sceptical prediction has a brave mate, called *credulous prediction*. A sentence is predictable in the credulous mode just in case it is contained in some particular extension of the premisses. (If there are no extensions, i.e. if the premisses are inconsistent, some stipulation has to be made.) Poole gives no hint as to how the choice is to be made; for the limited purpose of a purely formal study of credulous prediction we may assume that some *arbitrary* extension is chosen. Just as sceptical prediction corresponds to expectation inference based on full meet contractions, so does credulous prediction correspond to expectation inference based on maxichoice contractions.

On the assumption that the set D of defaults is closed under consequence, sceptical prediction is equivalent with expectation inference not only in the sense of (C'1) but also in the sense of (C1) (the definition favoured in [GM]). But this assumption leads to a serious degeneration of sceptical prediction. The exact correspondence between prediction according to Poole and full meet contraction allows to transfer results from the latter domain to the former. Among such transfers one, noted in Makinson (1994), makes use of Observation 2.15 to show that sceptical prediction turns so sceptical as to be useless if it is based on closed default sets: if D is a theory, then $C_D^{\text{SP}}(A) = \text{Cn}(A)$ in the principal case that D is inconsistent with A .

Merge inference

Default inferences according to Poole or the more general expectation-based inferences according to Gärdenfors and Makinson adjust hypotheses (or expectations) so as to suit the facts. But what, if the facts are unsuitable? That is to say, what, if the premisses for a defeasible inference are inconsistent?

Expectation-based inference—to focus on the most general notion of defeasible inference—models a way of reasoning in which one draws on powerful but risky hypotheses as much as one can without engendering

inconsistency. The reasoning proceeds on the basis of two collections of information: a set D of hypotheses or expectations (“defaults”) and a set A of “freshly supplied” facts or premisses. For the purpose of inference D and A are joined to a consistent whole. But consistency negotiations between A and D are not on equal terms: while A remains unchanged, D is being examined and pruned so as to obtain consistency with A .

The asymmetry of the situation has two unfortunate consequences. First, it places an absolute and supreme value on all members of a premise set A while allowing large variations in the value of members of the set D of background hypotheses. In principle, any member of D is vulnerable to rejection, while in any given inference from a set A , all members of A are given highest epistemic value. Many philosophers—most prominently Quine—have doubted whether hypotheses and evidence are ever used in this way in our efforts to obtain reliable information. In the passage cited in Chapter 1, James observes that confrontation with new evidence does not usually trigger a revision of one’s current beliefs so as to completely incorporate the new evidence. Indeed, the simplest and, as James remarks somewhat pessimistically, “by far the most usual” response to recalcitrant phenomena “is to ignore them altogether, or to abuse those who bear witness for them.” (1907, 61f.).

Second, inference that involves such asymmetry is only of limited use in averting the danger of inconsistency. For, while the set of hypotheses is always trimmed to consistency—even if it is inconsistent on its own—inconsistent premisses are never adjusted. Thus, if C is an expectation inference operation based on D , we will in general not have $D \subseteq C(A)$ but $A \subseteq C(A)$. Since expectation inference operations are *supraclassical*, i.e.

$$\text{Cn} \leq C$$

it follows that if A is inconsistent whence $\text{Cn}(A) = \text{Fml}$, then $C(A) = \text{Cn}(A) = \text{Fml}$.

Echoing Poole (1991, p. 27f.) there are two ways in which one may react to this finding.

- (1) The first is to claim that there is obviously something wrong with classical logic and so there is a need to define a new logic to handle reasoning from inconsistent premisses (e.g. da Costa (1974), Belnap (1977a,b)).
- (2) An alternative is to say that there is nothing wrong with classical logic; but we should not expect all reasoning to be straight classical or supraclassical deduction from our knowledge.

We shall now exhibit an approach to inference which, like Poole inference, is firmly based on classical logic—thus exemplifying the second

approach—but which is, unlike Poole inference, also suited for reasoning from inconsistent premisses.

All that is required for the purpose is to remove the imbalance in epistemic value between background hypotheses and explicit premisses. Instead of using a revision operation, which is biased towards the facts to be revised by, we shall use a *merge* operation which even-handedly negotiates between facts and hypotheses so as to obtain maximally reliable information.

Merge and paraconsistency. With the introduction of the merge operation we have just provided a tool for drawing non-trivial conclusions from inconsistent collections of premisses. We are thus in the vicinity of what has become known as *paraconsistent inference*. If we wish to emphasise the inferential aspect of the merge operation, we may define a relation \vdash of *merge inference* thus:

$$A \vdash \beta \iff A \circ A \vdash \beta$$

(In view of the discussion in the last section and in particular of the last Observation a number of equivalent definitions are possible.)

It is an immediate consequence of the definition that the relation \vdash does not enjoy the *overlap* property,

$$A \vdash \alpha, \text{ if } \alpha \in A$$

and that it does not even satisfy *reflexivity*,

$$\alpha \vdash \alpha$$

These facts may at first sight cast doubt on the presumption that \vdash is a relation of inference; for, reflexivity is usually seen as a fundamental condition on any *bona fide* inference relation. But there is another condition, equally fundamental, that, in the present context, conflicts with reflexivity and, hence, also with overlap. It is commonly agreed (see e.g. Makinson, 1994) that even the weakest defeasible inference relations should satisfy the condition of *right weakening*:

$$\frac{A \vdash \alpha_1 \quad \alpha_1 \vdash \alpha_2}{A \vdash \alpha_2}$$

Given reflexivity, we may always discharge the assumption $\alpha_1 \vdash \alpha_1$ and then derive from right weakening the condition of *supraclassicality*:

$$\frac{\overline{\alpha_1 \vdash \alpha_1} \quad \alpha_1 \vdash \alpha_2}{\alpha_1 \vdash \alpha_2}$$

But no supraclassical inference relation can serve the purpose of drawing non-trivial conclusions from inconsistent premisses. Thus, either reflexivity—or, by a similar argument, overlap—or right weakening has to be rejected. In the present context, the way the merge operation works motivates clearly a rejection of overlap and, hence, reflexivity. Note, however, that by Observation 4.12 merge inference is supraclassical with respect to consistent premisses; it therefore satisfies overlap whenever premisses need no adjustment.

The relation \vdash , as just defined, satisfies the only commonly agreed upon condition for paraconsistency: there are sets A and sentences β such that A is inconsistent and yet $A \not\vdash \beta$. Neither classical logic nor intuitionist logic are paraconsistent in this minimal sense. Although two important theories of inference are thus ruled out by the minimal condition, the condition lets pass muster many logics that would not count as paraconsistent in any substantial sense. For example, the condition is satisfied by the Kolmogorov-Johansson minimal logic. Yet minimal logic allows to infer arbitrary negated formulae from any inconsistent set of premisses. Though paraconsistent in letter minimal logic is thus not paraconsistent in spirit.

Urbas (199+) has shown that there are many—indeed infinitely many—ways in which an inconsistent set may inferentially be trivialised. The worst case of trivialisation occurs when every sentence of the language can be inferred. “Slightly” more discriminating are inference from inconsistency to every negation—as in minimal logic—or to every double negation, or to every conjunction in the language, and so on. It is impossible to find (finitely storable) conditions for paraconsistency that guard against all possible trivialisations.

Still, implicit in the work of most paraconsistent logicians there is a cluster of standards—of varying strength and sometimes pulling into opposite directions—as to what should be required of a substantive notion of paraconsistency. One of the pioneering researchers in paraconsistent logics, Newton da Costa, has summarised some of the standards guiding his own investigations as follows:

“I) In these [paraconsistent] calculi the principle of contradiction, $\neg(\alpha \wedge \neg\alpha)$, must not be a valid schema; II) From two contradictory formulas, α and $\neg\alpha$, it will not in general be possible to deduce an arbitrary formula β ; III) it must be simple to extend [these calculi] to corresponding predicate calculi (with or without equality) of first order; IV) [these calculi] must contain the most part of the schemata and rules of [classical logic] which do not interfere with the first conditions. (Evidently, the last two conditions are vague.)” (1974, p. 498)

It is remarkable that merge inference does very well with respect to almost all of these desiderata. Only the first condition is clearly not satisfied. For, on the understanding that a valid schema is one that \vdash -follows from the empty set, $\neg(\alpha \wedge \neg\alpha)$ is valid since we have $\emptyset \circ \emptyset \vdash \neg(\alpha \wedge \neg\alpha)$, i.e. $\vdash \neg(\alpha \wedge \neg\alpha)$. But as far as paraconsistent *inference* is concerned, condition I is not a particularly relevant integrity constraint as it pertains to theorems rather than consequences. Even when theorems are taken into view, the importance of the first condition can be disputed. The main concern about rejecting $\neg(\alpha \wedge \neg\alpha)$ as a theorem is that it comes at a heavy cost. For, given the usual DeMorgan, double-negation and substitution principles, $\neg(\alpha \wedge \neg\alpha)$ is equivalent to $\alpha \vee \neg\alpha$. Yet it does not seem essential to the idea of paraconsistency that one rejects the principle of the excluded middle. (Da Costa restricts the substitutivity of provable equivalents—an option which to many logicians appears to be the most awkward available. The option is, for example, heavily criticised by Routley *et al.* (1982).)

The remaining conditions are all satisfied by merge inference. Above we have already remarked on the second condition. As to the third condition, we just note that it is only for the sake of simplicity that we have taken into account only the sentential structure of language. Nothing of importance changes when richer structures are assumed.

Finally note that merge inference satisfies the fourth condition in da Costa's list even better than most of the usual paraconsistent logics. For, we can show (using Observation 4.12.3–4) that merge inference behaves entirely classically in the presence of consistent premisses. Moreover, if premisses are inconsistent, merge inference attempts to salvage a maximal amount of reliable information from the premisses by implementing a minimal mutilation policy.

There is no need to always resort to non-classical logics in order to solve the problem of how to draw non-trivial conclusions from inconsistent premisses. But merge inference is no universal tool for this task. Unlike paraconsistent logics—such as those from the family of relevant logics—merge inference is not suited to provide a notion of logical closure for many simple but inconsistent theories. In particular two such families of theories have driven much of the interest among logicians in paraconsistent logics: set theories with an unrestricted comprehension scheme, and truth theories with a naive T-scheme. Merge inference presupposes that maximally consistent subtheories of a given inconsistent theory can be determined and weighted. If either one of these presuppositions cannot be met, then only closure operations, as provided by paraconsistent logics, can avert the threat of triviality.

Despite *prima facie* similar concerns, research in paraconsistent logic on the one hand and the investigation of inconsistency handling in artificial intelligence on the other hand so far had very little impact on each other. An explanation of this state of affairs emerges from the last paragraph. Research in AI tends to be conservative whence fundamentally classically inspired. A typical technique—exemplified by the merge operation—for averting unfortunate side effects of the classical orientation consists in, first, pairing down inconsistent blocks of information into consistent sub-blocks, and, second, by adjudicating between rivalling consistent blocks of information by introducing some extra-logical structure (ordering, selection functions). There is thus a trade-off between the elegance and simplicity of the classical logical framework on the one hand, and the increased cost and limited feasibility of computing consistent sets and the injection of extra-logical information on the other hand. By and large, researchers in artificial intelligence work on the assumption that with respect to applications in artificial intelligence the balance tips in favour of the classical framework: that it is worth keeping classical logic as a basis and to pay for the repair kit.

Negation as cancellation. We close this chapter by pointing out a connection between the merge operation and its associated inference relation on the one hand and a perennial but—seen from a Boolean angle—heterodox view of negation on the other hand. The connection is yet another example of how topoi deeply rooted in the history of logic and sometimes worked out by formally inclined philosophers resurface in current research—despite the fact that in the past they have repeatedly been pronounced confused beyond recovery.

The view of negation we have in mind is sometimes used to motivate particular theories of paraconsistent consequence. According to this view, inconsistencies are, as it were, inferentially mute because their constituents “cancel” each other. Instead of inconsistencies having total content, they have no content at all; instead of everything following from a contradiction, nothing follows.

One of the advocates of the view in this century, P.F. Strawson (1952, pp. 2f.), compares someone who contradicts himself

“... with a man who makes as if to give something away and then takes it back again. He arouses expectations which he does not fulfill; ... The point is that the *standard* purpose of speech, the intention to communicate something, is frustrated by self-contradiction. Contradicting oneself is like writing something down and then erasing it, or

putting a line through it. A contradiction cancels itself and leaves nothing.”

This *cancellation* view of negation (the term was coined in Routley et al. 1982, pp. 89ff.; see also Routley and Routley 1983), though clearly at odds with the Boolean complementation picture, is a recurrent theme in the philosophy of logic. Berkeley, for example, advances the view in *The Analyst* (p. 73) as follows:

“Nothing is plainer than that no just conclusion can be directly drawn from two inconsistent premises. You may indeed suppose any thing possible: But afterwards you may not suppose anything that destroys what you first supposed: or, if you do, you must begin *de nove* ... [When] you ... destroy one supposition by another ... you may not retain the consequences, or any part of the consequences, of your first supposition so destroyed.”

It is worth observing that Berkeley clearly notes the non-monotonicity of any consequence relation based on a cancellation theory of negation.

In modern times there have been several attempts at casting the cancellation view of negation in a formal system—perhaps the most prominent being McCall’s (1967) system of ‘connexive implication’. All these systems constitute quite radical departures from classical logic. For example conjunction elimination is pronounced invalid because of instances such as

$$\alpha \wedge \neg\alpha \rightarrow \alpha$$

Some logicians, e.g. McCall, even advocate logical truths that are not classically recognised, such as Aristotle’s thesis—no proposition implies its own negation:

$$\neg(\alpha \rightarrow \neg\alpha)$$

For a critical survey of logics intended to formalise the cancellation view of negation the reader should consult Routley et al. (1982, Sec. 2.9).

A cancellation view of negation does not only find support in suggestive—but inconclusive—comparisons of negation with an operation of erasure (of content) but also in certain features of rational inquiry. In inquiry we seek information while trying to avoid error. These aims pull into opposite directions and therefore need balancing. In the presence of conflicting and equally trustworthy evidence, the danger of incurring error is particularly serious. Thus, when faced with inconsistent information it is frequently the best strategy to withhold judgement. As long as one does not know how a conflict is to be resolved, conflicting information is best be disregarded.

But to disregard conflicting information is simply to act on the assumption that no (reliable) information as to the truth-value of a proposition is, at present, available. Sometimes too much, i.e. conflicting information has the same value for us as too little, i.e. no information, namely none.

In merge inference we see a refined cancellation mechanism at work. To illustrate, consider a merge inference from a set $A = \{\alpha, \neg\alpha\}$. According to cancellation theories, inference from A should behave like inference from \emptyset . (It is a further and largely independent question as to what one should be able to infer from \emptyset .) This is, because α and $\neg\alpha$ cancel each other. Now, the recommendations of the cancellation and the merge theory coincide, if in the latter framework no distinction in value can be discerned between α and $\neg\alpha$. In that case merge inference will be of the full meet variety in which no selection is made among the maximally consistent subsets of a collection of premisses: they are all treated on a par.

In contrast to the recommendation of cancellation views, merge inference leaves room for the possibility that cancellation is asymmetric: that α overrides $\neg\alpha$ and not vice versa. Merge inference thus does justice to the valid intuition that sometimes conflicting evidence is better ignored. It does so, however, in a less dogmatic and more flexible manner. Perhaps most importantly, it shows how to accommodate this intuition in a well-understood classical framework, that is, without rushing into an under-motivated wholesale revision of logical theory.

Everything in Flux: Dynamic Ontologies

The reader who has followed the proofs recorded in previous chapters will not have failed to notice that they draw on very limited resources. Almost never do these proofs tap the internal structure of sentences; the only exceptions occur in connection with negation in Chapter Four. Apart from these sins provoked by the notion of revision, we never made any substantial assumptions about the relation of logical consequence; it was enough to know that logical consequence generates, in the familiar way, a closure operation.

Thus it appears that the notion of a contraction operations as developed in the previous chapters is not particularly tied to changes of theories or of beliefs. The theory of contractions allows for many more interpretations as it applies to a much wider class of phenomena than originally envisaged. The purpose of this final chapter is to give some flavour of just how general the theory really is. It also indicates some areas of application that seem particularly fertile ground for the theory thus pointing to promising future directions of research.

5.1 Closure systems and ontological spaces

Let \mathcal{F} be a family of subsets of some set X . The family \mathcal{F} is a *closure system* on X just in case it is closed under intersections and contains X :

$$\bigcap \mathcal{A} \in \mathcal{F}, \text{ for every nonempty } \mathcal{A} \subseteq \mathcal{F} \quad (\text{Intersection})$$

$$X \in \mathcal{F} \quad (\text{Top})$$

(*Notation:* Capital roman letters, A, B, C, \dots, X, Y, Z , (sometimes decorated) denote subsets of some domain U ; calligraphic capital letters, $\mathcal{A}, \mathcal{B}, \mathcal{C}, \dots, \mathcal{X}, \mathcal{Y}, \mathcal{Z}$, are reserved for families of subsets of U ; lower case roman letters, a, b, c, \dots , range variably over elements of the domain.)

An operation $\bar{}: \wp(X) \rightarrow \wp(X)$ is a *closure operation* on a set X , if it satisfies three conditions (for all $A, B \subseteq X$):

$$\begin{aligned} A &\subseteq \bar{A} && \text{(Reflexivity)} \\ \bar{A} &\subseteq \bar{B}, \text{ if } A \subseteq B && \text{(Monotonicity)} \\ \overline{\bar{A}} &\subseteq \bar{A} && \text{(Idempotence)} \end{aligned}$$

There is a well-known connection between closure systems and closure operations. (See e.g. Birkhoff 1967, Ch. V, Theorem 1, where closure systems appear under the name of *Moore families*.) Given a closure operation C on a set X , the family of closed subsets of X is a closure system: $\mathcal{X}_C = \{A \subseteq X: A = \bar{A}\}$. Conversely, given a closure system \mathcal{X} on X , the operation of intersecting all extensions in \mathcal{X} of some given set A is a closure operation on X : $C_{\mathcal{X}} = \bigcap\{B \in \mathcal{X}: A \subseteq B\}$. Moreover, the connection between closure operations and closure systems is a bijective one: $C_{(\mathcal{X}_C)} = C$ and $\mathcal{X}_{(C_{\mathcal{X}})} = \mathcal{X}$. It is therefore purely a matter of convenience whether a family of sets is studied *qua* closure system or *qua* system of closed sets. We shall here give preference to the latter option.

Remark: Abstraction could be carried further by (1) defining closure operations and closure systems on arbitrary partially ordered sets (posets) and (2) exploiting the bijective correspondence between a complete sublattice L of a poset and the set of fixpoints of the closure operation $\bar{a} = \bigwedge\{x \in L: x \leq a\}$. But, as yet, no useful purpose seems to be served by proceeding further in this direction. Let us note, however, that results concerning contractions continue to hold *modulo modulorum* if talk about sets with their natural order relation is replaced by talk about complete lattices. *End of remark*

Closure systems are everywhere: they arise naturally not only in all branches of mathematics but in just about every field of inquiry. To focus attention on philosophical matters, let us borrow the suggestive terminology of Fine (1991). Henceforth a closure system will be referred to as an *ontological space*. An *ontology* is a point in ontological space, i.e. a closed set.

This terminology can be justified as follows. Ontologies, in a pre-systematic sense, are *possible collections of objects* of a certain kind. To restrict attention to possible collections is a minimal attempt at shunning triviality.

Thus, if A is a collection of objects, then the collection consisting of A together with some object b may not be possible because some objects in A may exclude the presence of b . By the same token, A may not be possible without containing c because the presence of some objects in A may require the presence of c .

A relation of *requirement* arises in many contexts. For example, sets require their members, things require their parts, actions require preceding actions to establish necessary preconditions, the truth of certain sentences requires the truth of certain others, and so on. A collection of objects is in a natural sense incomplete and indeed impossible unless it contains all objects required for their joint existence. In other words: a collection of objects may fail to constitute a proper ontology unless it is closed under the requisite relation of requirement.

Requirement, in this general sense, has always captured the attention of philosophers. A particularly deep study of this notion can be found in the third of Husserl's *Logical Investigations*. For a more recent study of requirement relations see e.g. Fine (1991).

5.2 Travelling in ontological space

Moves in ontological space — transitions from one ontology to another — proceed by expanding or by shrinking or a by mixtures of both. Transitions by *expanding* pose no problem. To expand a given ontology A by some object b , it suffices to form the union of A and b and then to close: $\overline{A \cup \{b\}}$.

Shrinking, i.e. *contracting* an ontology can pose a problem. To illustrate, consider a “small” ontology

$$A = \overline{\{a_1, a_2\}}$$

such that a_1 and a_2 jointly but not separately require some object b :

$$b \in A \text{ but neither } b \in \bar{a}_1 \text{ nor } b \in \bar{a}_2 \quad (*)$$

(We abbreviate $\overline{\{a\}}$ to \bar{a} .) Which ontology should be the result of shrinking A just enough so as not to require b ? Should a_1 or a_2 or both be absent from the ontology that is the most like A except for not containing b ?

It was this sort of choice situation that led us (in Chapter Two) to adopt the partial meet modelling of contractions. Careful perusal of the results presented there reveals that this modelling transfers without any modification necessary to contractions defined on arbitrary closure systems, i.e. ontological spaces.

Given an ontological space, any mapping in that space is called a *transition function*. Two transition functions τ_1 and τ_2 are equivalent if $\tau_1(A) = \tau_2(A)$, for every ontology A .

The (unique) expansion operation together with a contraction operation satisfying the conditions in Chapter Three constitute a *functionally complete* set of transition functions for any ontological space. That is to say, for any transition function there exists an equivalent transition function which is a sequence of expansions and contractions. (In the “worst case” one can always contract to the empty ontology and then expand to the target ontology.) Since expansions pose no problem beyond that of specifying a suitable closure operation, a complete theory of transition functions, is essentially given by a theory of contractions only.

In general, there will be more than one transition function from a given ontology to some other ontology. Some of these equivalent transition functions will be more direct or efficient than others; others will be more round-aboutish. This opens the prospect of studying the domain of transition functions, say, from the viewpoint of computational complexity — a prospect of obvious practical importance.

5.3 Finding (good) examples

Closure systems are everywhere; and our theory of contractions covers all changes within arbitrary closure systems. These findings harbour the suspicion that the theory of contractions is in danger of passing a threshold of generality beyond which triviality lurks. Thus, every domain of ontologies that can be generated by some closure operation will serve as an example for the theory. It seems that there is little left that cannot be bent to exemplify the theory in some way or another. We shall now attempt to answer this suspicion by distinguishing between good examples or applications and bad or mostly uninteresting ones.

Describing examples of ontologies with contraction functions requires at least identifying

- (I) a *domain* of objects to make up ontologies, and
- (II) an appropriate relation of *requirement*.

These two steps set up an ontological space, as explained above, in which contraction functions are well-defined.

However, contraction functions lose much of their interest if they are applied to domains in which no choice can arise or in which no differences of value are discernible. Such domains give rise to *bad* examples of our theory of contractions.

Recall that we have motivated our abstract theory of contractions by considering examples of the kind where some set, say $\{a, b\}$, requires an element c without it being the case that the requirement can be traced down to any single elements of that set; thus, neither $\{a\}$ nor $\{b\}$ may require c . Obviously, such examples only arise if the underlying closure operation violates the condition of “maximal compactness”:

$$\text{if } b \in \overline{A}, \text{ then } \exists a \in A \text{ such that } b \in \overline{a}$$

or, writing $A \longrightarrow b$ for $b \in \overline{A}$:

$$\text{if } A \longrightarrow b, \text{ then } a \longrightarrow b \text{ for some } a \in A \quad (\text{Maximal Compactness})$$

(Here and in the rest of this paper we revert to the conceptually simpler notion of contracting by a single object. The above condition and the following results are easily extended to multiple contractions.) It can be shown that if the closure operation is maximally compact then (partial meet) contraction for singletons is definable—like expansion—purely in terms of set-operations and closure. Define

$$\downarrow a = \{x \in D: x \longrightarrow a\}$$

where D is the domain of objects under consideration.

OBSERVATION 5.1 *If the closure operation $\overline{(\quad)}$ is maximally compact, then $A \perp b = A \setminus \downarrow b$, for all ontologies A and objects b .*

PROOF. It will suffice to show that $A \perp b = \{A \setminus \downarrow b\}$. For, then it follows that

$$A \setminus \downarrow b = \bigcap s(A \perp b) = A - b$$

as required.

Claim 1: $A \setminus \downarrow b \in A \perp b$.

Clearly, (a) $A \setminus \downarrow b \subseteq A$. Also, (b) $A \setminus \downarrow b \not\rightarrow b$. For, otherwise by maximal compactness there would be an $x \in A \setminus \downarrow b$ such that $x \longrightarrow b$, i.e. $x \in \downarrow b$ —contradiction. Finally, (c) suppose there is some set B such that $A \setminus \downarrow b \subset B \subseteq A$ and $B \not\rightarrow b$. Then the first conjunct would require that $B \cap \downarrow b \neq \emptyset$ whereas the second would require the contradictory. Thus we have verified the three conditions for membership in the set $A \perp b$ of remainders.

Claim 2: If $B \in A \perp b$ then $B = A \setminus \downarrow b$. Suppose that $x \notin A \setminus \downarrow b$. Then either (a) $x \notin A$ or (b) $x \in \downarrow b$. If (a), then $x \notin B$ since $B \subseteq A$. If (b), then $x \longrightarrow b$. Hence, $x \notin B$ since $B \not\rightarrow b$. So $x \notin B$; hence, $B \subseteq A \setminus \downarrow b$. Since both B and $A \perp b$ are assumed to be in $A \perp b$, one cannot be a proper subsets of the other. Thus $B = A \setminus \downarrow b$. ■

There was no need to guard against maximal compactness when theory change was the only application in view. Logical consequence is (almost) never maximally compact. (The only exception appears to be Jaśkowski's relation of discussive consequence: see Priest and Routley (1983), 111.) But now, with a wider perspective, maximal compactness is sometimes a most natural property. A man with a heart disease may need a pacemaker. He may be a member of a group of patients. So there is a natural sense in which that group of patients needs a pacemaker and, perhaps, also a kidney. However, the group does not need these things collectively but individually—in that sense, the needs of the group are maximally compact.

In ontological spaces with a maximally compact closure operation no choice situation can arise when trying to contract; hence, contractions are completely determined by the familiar set operations. To find an interesting application for our theory the ontological space will have to be generated by

(III) a closure operation that is *not maximally compact*.

A second source of trivialisation is an egalitarian domain. If the domain carries no value structure to make nonarbitrary selections among collections of objects, then the selection function central to the partial meet approach degenerates to identity. In that case we have

$$A - B = \bigcap s_A(A \perp B) = \bigcap (A \perp B) \quad (*)$$

for all subsets A and B of the domain. This makes $A - B$ a very small subontology of A — too small to pass the condition that contractions should perform minimal *incisions* into an ontology. (Cf. AGM's observations on "full meet" contractions cited in Observation 2.15.) Thus there is a fourth condition for identifying good examples, viz. that the ontological space under consideration should be based on

(IV) a domain with a *value structure*.

Note that, in contrast to the notion of closure under what is required, there is no need to make that structure explicit — it suffices to assume the existence of such a structure.

Thus guarded against bad examples of ontological spaces with contraction functions let us now turn to good ones apart from belief sets and theories.

Actions (1). Consider a domain D of actions. Some actions require that certain others be carried out. For example, before potatoes can be boiled they need to be peeled or scraped. Thus, an action is *performable* only

if all actions it requires are performable; and a set of actions is performable only if each of its elements are performable. The relation

... makes ... performable

generates a closure operation. The ensuing ontological space is the collection of all performable sets of actions, taken from the domain D .

The *makes-performable* relation can be “massively disjunctive” because both a and a' may make b performable without b requiring any particular one of a or a' . (Before our heart patient can undergo surgery a pacemaker—but not a particular one!—needs to be supplied.) Do we want both a and a' in the closure of b ? There are a number of strategies to prevent ontological inflation at this point. One is to complicate the definition of closure from the *makes-performable* relation. Another is to introduce more structure to the domain such that the notion of a “disjunctive” action becomes expressible (as in the language of dynamic logic). A third strategy may be to move from a domain of action tokens to one of judiciously chosen action types.

To satisfy condition (III), we need to look for situations such that if two actions, a and b , are *both* to be performed, then some action c is required, while c may not be required if only one of a or b is to be performed. Hence, if for some reason c cannot be carried out, then not both a and b but, perhaps, one of the two can be performed — a choice has to be made.

Many situations of this kind are of the *scarce-resources* type. The joint performance of a and b may require resources, to be provided by some action c , which are not required if only the one or the other action is to be carried out.

Another class of situations leading to choice is of the *coordination* type. Doing both a and b may require some coordinating action c which would not be required if only the one or the other of the two actions were performed.

Actions (2). There are other requirement relations definable on a domain of actions that lead to different ontological spaces. Think, for example of a domain of programs and let the requirement relation be specified by the set of call procedures,

program ... calls program ...

Again, this relation induces a closure operation and, hence, an ontological space. Choice can arise, as before, in situations that require coordinating. For example, a and b may be programs which, if run in parallel, call a routine c that prevents interference. If routine c cannot be run, only one of a or b can be executed at a time.

Things. Things require their parts. Already on the very weak assumption that *some* things may fuse with others, maximal compactness fails and so contraction may demand choice. For, if for some individuals a and b , the fusion ab exists as soon as a and b exist, then the individuals a and b collectively require the fusion ab without a or b individually requiring ab . Hence, a fusion can only be withdrawn if at least one of its constituents is withdrawn. (When mereology is put to work one usually assumes a much stronger fusion principle, requiring that any collection of individuals fuses; cf. Lewis (1991), sec. 3.4.)

Comparative necessity. In each ontological space a class of ubiquitous, or necessary objects is defined: those objects are necessary that are contained in each ontology.

Necessary objects cannot be deleted from any ontology; other objects can. But not all objects are as easily deleted as any others. In situations where a choice is required some objects may exhibit more resistance to deletion than others. If we identify absolute undeletability with necessity, then comparisons of deletability may be taken as comparisons of necessity. Some objects give way more easily than others; hence, they are less necessary, their absence in an ontology is more possible.

Let ab be a “strongest” object required by the set $\{a, b\}$. By this we mean: $a, b \longrightarrow ab$ and for no c does $a, b \longrightarrow c \longrightarrow ab$ while $ab \not\rightarrow c$. (Strongest objects are unique up to \longleftrightarrow .) Then we may define a to be no more necessary than b ($a \leq b$) (given some fixed ontological space \mathcal{S}) just in case b (but perhaps not a) always survives contraction by ab . More formally:

$$a \leq b \text{ iff } \forall X \in \mathcal{S}: b \in X - ab \quad (\text{CN})$$

(Compare the relation \leq with that of epistemic entrenchment, as explained in Gärdenfors 1988.) Clearly, under the definition (CN) ubiquitous elements carry maximal necessity. Properties of comparative necessity may now be studied vis-à-vis conditions on contraction functions.

Functional dependencies. Finally, consider a domain of *attributes* mapping *objects* to certain *values*. For example, the attribute ‘temperature’ would map an object (of the right kind) to a specific value on a suitable scale, the attribute ‘colour’ would map an object to a specific shade, the attribute ‘boils’ would map an object to one of the two values ‘yes’ or ‘no’, and so on. Information about particular objects may be represented as a tuples with each position in the tuple reserved for a value for a particular attribute. For example the tuple

$$\langle \text{H}_2\text{O}, 120, Y \rangle$$

may be used to represent the information that a particular chunk of matter is made of water (H_2O), that it has a temperature of 120 degrees Celsius and that it boils (Y). Obviously, this information contains an element of redundancy since water always boils at 120 degrees Celsius. The value of ‘boils’ depends on the kind of material at issue and its mean temperature. That is to say, for every two objects x and y , if x and y are alike in material and temperature, then they agree on their boiling point:

$$\text{mat}(x) = \text{mat}(y) \ \& \ \text{temp}(x) = \text{temp}(y) \quad \implies \quad \text{boil}(x) = \text{boil}(y)$$

where the individual variables x and y are universally quantified over a fixed domain of objects. Such statements of functional dependency for a given domain of attributes, objects and values may be shortened to

$$\text{mat}, \text{temp} \longrightarrow \text{boil}$$

The notion of functional dependencies is obviously of practical importance in the design of good, i.e. irredundant databases. Under the name of ‘supervenience’ it is also a much discussed notion in philosophy, particularly in metaphysics.

The key observation for connecting functional dependencies with contraction operations has already been recorded in Chapter Two (page 15): the relation of Bolzano consequence—used for defining theories and package contractions—has the same formal properties as the relation of functional dependency as characterised by the Armstrong axioms (Armstrong, 1974). Thus, a set of attributes closed under functional dependency constitutes a closure system, an ontology of properties. Given an object x and a set of attributes A that are assumed to have determinate values for x , closure under functional dependency generates the set \overline{A} of all attributes that must also be determinate for x . If determination is undone for any particular property in \overline{A} , a contraction procedure will have to determine which other properties in \overline{A} can no longer be assumed fixed.

There is more work for contraction here. The transitive closure of a set of functional dependencies is itself a closure system. It is sensitive to the addition of new entries in a database: as new information is added closure under functional dependency may engender inconsistency. In such a situation one may not always want to retract data. After all the database may have the status of a collection of “hard” facts whereas the dependencies may be no more than rules of thumb. If this is the case, then conjectured dependencies have to respect the facts and it is the former that will undergo contraction.

References

- ALCHOURRÓN, C.E., P. GÄRDENFORS and D. MAKINSON, On the logic of theory change: partial meet contraction and revision functions, *Journal of Symbolic Logic* 50 (1985), 510–530.
- ALCHOURRÓN, C.E. and D. MAKINSON, Hierarchies of regulations and their logic, in *New Studies in Deontic Logic*, ed. by R. Hilpinen, Dordrecht (Reidel), 1981, 125–148.
- ALCHOURRÓN, C.E. and D. MAKINSON, On the logic of theory change: Contractions functions and their associated revision functions, *Theoria* (1982), 14–37.
- ANDERSON, A.R. and N.D. BELNAP, *Entailment. The Logic of Relevance and Necessity*, 1, Princeton (University Press), 1975.
- ARMSTRONG, W.W., Dependency structures of database relationships, *Proc. 1974 IFIP Congress*, Amsterdam (North-Holland), 1974, 580–583.
- ARRUDA, A.I., R. CHUAQUI and N.C.A. DA COSTA (EDS.), *Mathematical Logic in Latin America*, Amsterdam (North-Holland), 1980.
- BACH, K., Default reasoning: Jumping to conclusions and knowing when to think twice, *Pacific Philosophical Quarterly* 65 (1984), 37–58.
- BECK, S., L. BERTHOLLE and J. CHILD, *Mastering the Art of French Cooking*, vol. 1, Harmondsworth (Penguin), 1966.
- BELNAP, N.D., How a computer should think, in *Contemporary Aspects of Logic*, ed. G. Ryle, Stocksfield (Oriel Press), 1977a.
- BELNAP, N.D., A useful four-valued logic, in *Modern Uses of Multiple-Valued Logic*, ed. J.M. Dunn and G. Epstein, Dordrecht (Reidel), 1977b.
- BERKELY, G., *The Analyst*, in *The Works of George Berkeley*, ed. A.A. Luce and T.E. Jessop, vol. 4, London (Nelson & Sons), 1951.
- BIRKHOFF, G., *Lattice Theory*, Providence (AMS), 3rd ed., 1967.
- BOLZANO, B., *Wissenschaftslehre*, 1837.
- BREWKA, G., *Nonmonotonic Reasoning: Logical Foundations of Commonsense*, Cambridge (University Press), 1991.
- DA COSTA, N.C.A., On the theory of inconsistent formal systems, *Notre Dame Journal of Formal Logic* 15 (1974), 497–510.
- DUBOIS, D. and H. PRADE, Belief change and possibility theory, in *Belief Revision*, ed. P. Gärdenfors, Cambridge (University Press), 1992.

- FAGIN, R. and M.V. VARDI, The theory of data dependencies — an overview, in *Automata, Languages and Programming, LNCS 174*, ed. J. Paredaens, Berlin (Springer), 1984.
- FINE, K., The study of ontology, *Noûs* 25 (1991), 263–294.
- FUHRMANN, A., *Relevant Logics, Modal Logics, and Theory Change*, PhD thesis, Australian National University, Canberra, 1988.
- FUHRMANN, A., Reflective modalities and theory change, *Synthese* 81 (1989), 115–134.
- FUHRMANN, A., Theory contraction through base contraction, *Journal of Philosophical Logic* 20 (1991), 175–203.
- FUHRMANN, A., Travels in ontological space, *Proceedings of Analyomen 2*, ed. G. Meggle and P. Steinacker, Berlin–New York (de Gruyter), 1995.
- FUHRMANN, A., Everything in flux: Dynamic ontologies, forthcoming, 199+a.
- FUHRMANN, A., Nonmonotonic reasoning, in *Routledge Encyclopedia of Philosophy, Logic and Philosophy of Mathematics*, ed. M. Detlefsen, London (Routledge), 199+b.
- FUHRMANN, A. and S.O. HANSSON, A survey of multiple contractions, *Journal of Logic, Language and Information* 3 (1994), 39–76.
- FUHRMANN, A. and I. LEVI, Undercutting and the Ramsey test for conditionals, *Synthese* November (1994); also in *Learning Dynamics and Norms*, ed. C. Bicchieri, R. Jeffrey and B. Skyrms, Cambridge (University Press), 1995.
- FUHRMANN, A. and H. ROTT (eds.), *Logic, Action, and Information. Essays on Logic in Philosophy and Artificial Intelligence*, Berlin - New York (de Gruyter), 1996.
- GABBAY, D.M., C.J. HOGGER and J.A. ROBINSON, *Handbook of Logic in Artificial Intelligence and Logic Programming*, v.3: Nonmonotonic Reasoning and Uncertain Reasoning, Oxford (University Press), 1994.
- GÄRDENFORS, P., Rules for rational changes of belief, in *(320311): Philosophical Essays Dedicated to Lennart Åqvist on his Fiftieth Birthday*, ed. T. Pauli, Uppsala (Uppsala University), 1982.
- GÄRDENFORS, P., *Knowledge in Flux*, Cambridge, Mass. (MIT Press), 1988.
- GÄRDENFORS, P., Belief revision and nonmonotonic reasoning: Two sides of the same coin?, in *ECAI 90*, ed. Aiello, London (Pitman), 1990.
- GÄRDENFORS, P. and D. MAKINSON, Nonmonotonic inference based on expectations, *Artificial Intelligence* 65 (1994), 197–245.
- GÄRDENFORS, P. AND H. ROTT, Belief revision, in *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. 4: Epistemic and Temporal Reasoning, ed. D.M. Gabbay et al., Oxford (Clarendon Press), 1995, 35–132.
- GENTZEN, G., Untersuchungen über das logische Schliessen, *Mathematische Zeitschrift* 39 (1934–35), 176–210, 405–431.

- GINSBERG, M.L., Introduction, in *Readings in Nonmonotonic Reasoning*, ed. M. Ginsberg, Los Altos (Morgan Kaufmann), 1987; also as: AI and nonmonotonic reasoning, in *Handbook of Logic in Artificial Intelligence and Logic Programming*, 3, ed. D.M. Gabbay et al., Oxford (University Press), 1994.
- GROVE, A., Two modellings for theory change, *Journal of Philosophical Logic* 17 (1988), 157–170.
- HANSSON, S.O., New operators for theory change, *Theoria* 55(1989), 114–132.
- HANSSON, S.O., Belief revision without recovery, *Studia Logica* 50(1991), 251–260.
- HANSSON, S.O., Non-prioritized belief change, in *Belief Base Dynamics*, Uppsala University, 1991.
- HANSSON, S.O., A dyadic representation of belief, in *Belief Revision*, ed. by P. Gärdenfors, Cambridge (University Press), 1992a, 89–121.
- HANSSON, S.O., Similarity Semantics and Minimal Changes of Belief, *Erkenntnis* 37 (1992b), 401–429.
- HANSSON, S.O., Reversing the Levi identity, *Journal of Philosophical Logic*, 22 (1993), 637–669.
- HUDSON, J.L., Logical subtraction, *Analysis* 35 (1975), 130–135.
- HUMBERSTONE, L., Logical subtraction: problems and prospects, typescript, Monash University, Melbourne, 1981.
- HUSSERL, E., *Logische Untersuchungen*, 2nd rev. ed., 1913.
- JAEGER, R.A., Action and subtraction, *Philosophical Review* 82 (1973), 320–329.
- JAMES, W., *Pragmatism—A New Name for Some Old Ways of Thinking*, New York (Longmans, Green, & Co.), 1907.
- JAY, R., *Learned Pigs and Fireproof Women*, New York (Warner Books), 1986
- KATSUNO, H. and A.O. MENDELZON, Propositional knowledge base revision and minimal change, *Artificial Intelligence* 53 (1991), 263–294.
- KATSUNO, H. and A.O. MENDELZON, On the difference between updating a knowledge base and revising it, in *Belief Revision*, ed. P. Gärdenfors, Cambridge (University Press), 1992.
- KONOLIGE, K., Abduction versus closure in causal theories, *Artificial Intelligence* 53 (1992), 255–272.
- KUHN, T.S., *The Structure of Scientific Revolutions*, Chicago (Chicago University Press), 1962.
- LEVI, I., *Gambling with Truth*, New York (A. Knopf), 1967.
- LEVI, I., Subjunctives, dispositions and chances, *Synthese* 34 (1977), 423–455.
- LEVI, I., *The Enterprise of Knowledge*, Cambridge, Mass. (MIT Press), 1980.

- LEVI, I., *The Fixation of Belief and Its Undoing*, Cambridge (University Press), 1991.
- LEVI, I., *For the Sake of the Argument. Ramsey Test Conditionals, Inductive Inference, and Nonmonotonic Reasoning*, Cambridge (University Press), 1996.
- LEWIS, D., *Counterfactuals*, Oxford (Blackwell), 1973.
- LEWIS, D., *Parts of Classes*, Oxford (Blackwell), 1991.
- LINDSTRÖM, S., A semantic approach to nonmonotonic reasoning, *Uppsala Preprints in Philosophy*, no. 6, Department of Philosophy, Uppsala University, 1991.
- LINDSTRÖM, S. and W. RABINOWICZ, Epistemic entrenchment with incomparabilities and relational belief revision, in *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence, No. 464, ed. A. Fuhrmann and M. Morreau, Berlin - Heidelberg - New York (Springer), 1991, 93–126.
- LYCAN, W.G., *Judgement and Justification*, Cambridge (University Press), 1988.
- MCCALL, S., Connexive implication, *Journal of Symbolic Logic* 31 (1966), 415–453.
- MAKINSON, D., How to give it up, *Synthese* 62 (1985), 347–363.
- MAKINSON, D., On the status of the postulate of recovery in the logic of theory change, *Journal of Philosophical Logic* 16 (1987), 383–394.
- MAKINSON, D., General Patterns in Nonmonotonic Reasoning, in *Handbook of Logic in Artificial Intelligence and Logic Programming*, 3, ed. D. Gabbay et al., Oxford (University Press), 1994.
- MAKINSON, D. and P. GÄRDENFORS, Revisions of knowledge systems using epistemic entrenchment, in *TARK II*, ed. M. Vardi, Los Altos (Morgan Kaufmann), 1988.
- MAKINSON, D. and P. GÄRDENFORS, Relations between the logic of theory change and nonmonotonic logic, in *The Logic of Theory Change*, ed. by A. Fuhrmann and M. Morreau, LNAI 465, Berlin etc. (Springer), 1991, 185–205.
- MAKINSON, D. and P. GÄRDENFORS, Nonmonotonic reasoning based on expectations, *Artificial Intelligence* 65 (1994), 197–245.
- MEYER, R.K., *A Theory of Theories*, undated typescript, 164 pp., Logic Group, Australian National University, Canberra, 197+.
- NIEDERÉE, R., Multiple contraction, in *The Logic of Theory Change*, ed. by A. Fuhrmann and M. Morreau, LNAI 465, Berlin etc. (Springer), 1991, 322–334.
- PEIRCE, C.S., The Fixation of Belief (1877), in *Writings of Charles S. Peirce*, 1, ed. M. Fisch, Bloomington (University of Indiana Press), 1982.
- PITT, J.C. (ed.), *Theories of Explanation*, New York - Oxford (Oxford University Press), 1988.

- POOLE, D., A logic for default reasoning, *Artificial Intelligence* 13 (1988), 27–47.
- POOLE, D., Default logic, in *Handbook of Logic in Artificial Intelligence and Logic Programming*, 3, ed. D. Gabbay et al., Oxford (University Press), 1994.
- PRIEST, G. and R. ROUTLEY, *On Paraconsistency*, Research Series in Logic and Metaphysics, ANU, Canberra 1983.
- PRIEST, G. and R. SYLVAN (eds.), *On Paraconsistency*, Munich (Philosophia), 1989.
- QUINE, W.V.O. and J.S. ULLIAN, *The Web of Belief*, New York (Random House), 1970 (rev'd 1978).
- REITER, R., A logic for default reasoning, *Artificial Intelligence* 13 (1980), 81–132.
- ROTT, H., Two methods of constructing contractions and revisions of knowledge systems, *Journal of Philosophical Logic* 20 (1991),
- ROTT, H., Modellings for belief change: base contractions, multiple contractions, and epistemic entrenchment, *Konstanzer Berichte*, No. 20, March 1992a.
- ROTT, H., On the logic of theory change: More maps between different kinds of contraction functions, in *Belief Revision*, ed. P. Gärdenfors, Cambridge (University Press), 1992b.
- ROTT, H., Belief contraction in the context of the general theory of rational choice, *Journal of Symbolic Logic*, 58 (1993), 1426–1450.
- ROUTLEY, R. et al., *Relevant Logics and Their Rivals*, Atascadero (Ridgeview), 1982.
- ROUTLEY, R. and V. ROUTLEY, Negation and contradiction, in *Proceedings of the Fifth Latin American Symposium on Mathematical Logic*, New York (Marcel Dekker), 1983.
- SCOTT, D.S., Completeness and axiomatizability in many-valued logic., in *Proceedings of the Tarski Symposium*, ed. L. Henkin et al., Providence, Rhode Island (Amer. Math. Soc.), 1974.
- SEGERBERG, K., *Classical Propositional Operators*, Oxford (Clarendon Press), 1982.
- SHACKLE, G.L.S., *Decision, Order and Time in Human Affairs*, Cambridge (University Press), 1961 (2nd ed. 1969).
- SHOESMITH, D.J. and T.J. SMILEY, *Multiple-Conclusion Logic*, Cambridge (University Press), 1978.
- SHOHAM, Y., *Reasoning About Change*, Cambridge, Mass. (MIT Press), 1988.
- SIMONS, P., *Parts*, Oxford (O.U.P.), 1987.
- SPOHN, W., Ordinal conditional functions: A dynamic theory of epistemic states, in *Causation in Decision, Belief Change, and Statistics*, 2, ed. W.L. Harper and B. Skyrms, Dordrecht (Reidel), 1988.

- STRAWSON, P.F., *Introduction to Logical Theory*, London (Methuen), 1952.
- TARSKI, A., Fundamental concepts of the methodology of the deductive sciences (1930), in *Logic, Semantics, Metamathematics*, ed. J.H. Woodger, Oxford (Clarendon), 1956.
- ULLMAN, J.D., *Principles of Database and Knowledge-Base Systems*, vol. 1, Rockville (Computer Science Press), 1988.
- ULLMANN-MARGALIT, E., On presumptions, *The Journal of Philosophy* 80 (1973), 143–163.
- URBAS, I., On paraconsistency, 199+.
- WITTGENSTEIN, L., *Philosophical Investigations/Philosophische Untersuchungen*, Oxford (Blackwell), 1953.
- WÓJCICKI, R., *Theory of Logical Calculi*, Dordrecht (Kluwer), 1988.