

Information structures in belief revision

Hans Rott

30 August 2006, 21:45

1 Introduction

Belief revision theories address the problem of rationally integrating new pieces of information into an agent's belief state. Belief states themselves are represented by certain kinds of information structures. The most interesting and most closely studied type of belief revision is the one in which the new information is inconsistent with the agent's current beliefs. The main question of the present contribution is: What kinds of information structures have figured prominently in belief revision theories as they have been developed over the past 25 years?

My focus will be on the philosophy behind the 'classical belief revision', that is of the AGM paradigm (so-called after its founders Alchourrón, Gärdenfors and Makinson and their 1985 paper) and its later generalizations. I shall talk about the generalizations to iterations of belief change in the 1990s and to operations of belief fusion or merging in the present century.

I am going to presuppose that information comes in symbolic form. There is no discussion of signals, symptoms, pictures, and similar non-codified formats in which information could be embodied. I shall say very little about approaches that make essential use of numerical, quantitative information. This of course is not to suggest that probabilistic approaches, evidence theory (Shafer 1976) or ranking functions (Spohn 1988, Goldszmidt and Pearl 1992) are not interesting and useful. It is just that qualitative approaches using preferential structures have played a major role in belief change theories, and they easily offer enough material for a chapter of their own.

The present chapter should be read in conjunction with the contribution by Baltag, van Ditmarsch and Moss (2006). In several respects, their approach is wider than the one presented here. However, the particular emphasis of belief revision theory is on the case where new information conflicts logically with the information previously accepted (in the terminology to be introduced later: with the old 'data base'). It can thus be seen as providing a module that can be plugged in, as it were, into the framework of Baltag et al. The following

presentation, however, is in the style of the more traditional, classical work in the area.¹

The plan of this chapter is as follows. Section 2 contains some preliminary remarks on information and truth, and on belief change as embedded in a functionalist philosophy of mind. Section 3 presents the problem of belief change as being a compound of a processes of reflection and processes of revision. Depending on which of these processes are given center stage, we can distinguish foundationalist and coherentist approaches to belief change (in roughly the sense that is known from contemporary epistemology). Section 4 gives three simple examples of such approaches. There are approaches that look coherentist but can be reconstructed as generated through hidden foundationalist recipes. Reconstructions of such a flavour are fairly typical of belief revision theories. Section 5 traces the idea that the static picture of belief as being represented by an information structure may encode much of the dynamics of belief. I shall interpret the history of belief revision theory of the last three decades as a story of finding an appropriate notion of a belief state, and its interpretation as providing a framework for analyzing both static and dynamic aspects of belief.

2 Preliminary remarks on information, truth and mind

2.1 Remarks on information and truth

There are countless explications of the term ‘information.’ I would like to propose the following informal and very general definition:

Information is some structure realized in the physical world that is suitable to be interpreted or exploited by some receiver in a reasonable way.

By *interpretation* or *exploitation*, I mean some kind of causal interaction between two entities, not necessarily human or even living, with a certain asymmetry between input and receiver. Interpretation in this sense is extremely wide (too wide perhaps), it may possibly, but need not necessarily make use of cognitive or linguistic means. A *reasonable* way is one that ‘makes sense’ of the information structure, and leads to ‘successful’ behaviour or action of the interpreter.

¹For important attempts to transfer carefully belief revision theories into the framework of modal logic, see Fuhrmann (1991), Cantwell (1997), Lindström and Rabinowicz (1999), Segerberg (2001) and van Benthem (to appear). For an enlightening discussion of the merits of the somewhat non-standard attitude towards logic in the AGM program, see Makinson (2003).

Often information, or rather pieces of information, are taken to represent states of affairs or objects. Some kinds of information (*signals*) are *truthful* (or *veridical*), but they may still be *deceptive* in the sense that they give rise to an inadequate interpretation, e.g., to false conclusions or unsuccessful behaviour. In such cases, if there is no natural or necessary link between a piece of information and what it represents, i.e., if it is *symbolic* and has a conventional meaning, then information may itself be called *false* (or *misinformation* or *misrepresentation*). However, it does not seem to be a fact of the matter whether the ‘falsity’ of the link between information and what it represents is ‘due to’ the carrier of information itself or the receiver interpreting this unit of information.

What is information as it figures in theories of belief revision? *Information* is that which may enter some belief state and change or transform it into another belief state. Belief revision theories usually do not care about the truth of beliefs – nor do they address the question of their justification or reliability.² For this reason, widely used terms and phrases like ‘knowledge base’, ‘knowledge representation’ and ‘knowledge in flux’, though widely used, are misnomers.

The idea, also emphasized in dynamic semantics or update semantics,³ is to characterize (the content of) a piece of information by the transformations of the receivers’ internal states that it brings about:

$$\langle \text{prior-state} \rangle \mapsto \langle \text{posterior-state} \rangle$$

One can say that such a transformation captures the interpretation or exploitation of the piece of information mentioned above.

Information in belief revision is

- *syntactic* in the sense that it is representable by sentences of some appropriate systematic language, and that it can be combined in the typical way sentences can,
- *semantic* in the sense that
 - it is not sensitive to transformations into logical equivalents,⁴ and
 - as long as the new information is consistent with the current belief state, it simply rules out possibilities.⁵

²But see Kelly, Schulte and Hendricks (1997) and Kelly (1999).

³See Stalnaker (1984), Gärdenfors (1988, chapter 6), Groenendijk and Stokhof (1991), and Veltman (1996)

⁴Compare, e.g., axioms (AGM6) and (DP6) below.

⁵Compare, e.g., axioms (AGM3), (AGM4), (DP3) and (DP4) below.

2.2 Some clues from the philosophy of mind

In belief revision, the main interest is in information for receivers with a mind, i.e., humans.⁶ The philosophy of mind is concerned with the relation between the mind and the body of human beings (compare Kim 2006, chapters 2–6). How can it be that a person, that is first of all a biological organism, exhibits thoughts, desires, feelings, etc? How can these mental states, states that seem to be perfectly accessible to the persons that are in them but not to any third persons, be the objects of scientific studies? Let us call the mental life of a person his or her *psychology*.

According to behaviourism, the leading school of psychology until the middle of the 20th century, a person's psychology can be identified with (characterized by, reduced to) the function

$$\text{input} \longmapsto \text{output}$$

where inputs and outputs, (physical) stimuli and (behavioural) responses, are observable entities. Behaviourism turned out to be too simplistic. Another approach to objectify the human mind was provided by the physicalist or materialist identity theory of mind. According to this approach, *a person's psychology* can be identified with (characterized by, reduced to) its physical or material state. But this idea would rule out beings with a different physiology like non-human animals or Martians or machines from having psychological states, something one would at least want to leave room for.

So a third paradigm, functionalism, appeared on the scene which in a way combined the best of the previous approaches. Functionalism is behaviorism plus internal states, or – approaching it from the other side – functionalism is materialism plus multiple realizability.⁷ Alan Turing (1950) and Hilary Putnam (1960, 1967) promoted the idea that human thinking could be likened to the calculations that go on in a computer (a Turing machine). The *computer metaphor* became popular which says that mind is to brain as software is to hardware. A software is a program that can be described abstractly by a (finite) machine table specifying completely all (of a man's or a computer's) transitions of the following type:

$$\langle \text{prior-state, input} \rangle \longmapsto \langle \text{posterior-state, output} \rangle$$

Prior and posterior states are *internal states* or *psychological states* of the person or the computer. The set of all such transitions was called a 'machine table' by the early functionalists, and it fully specifies *a person's psychology* or a *computer program*.

⁶Or information for computing machinery, but we will see that this difference is not important for our purposes.

⁷The thesis of multiple realizability says that one mental state can be 'realized' or 'implemented' by varying physical states. Beings with different physical constitutions can thus be in the same mental state.

2.3 Functionalism as applied to belief revision

Belief is an internal affair. Any possible announcement of one's beliefs or any other action following a belief may be disregarded for the purposes of this paper. Thus, we do not need to deal with any manifest output.⁸ This suggests that the functionalist format can be reduced to a simpler mapping of the following format:

$$\langle \text{prior-state, input} \rangle \mapsto \langle \text{posterior-state} \rangle$$

The *input* is a piece of information. This actually marks the proper place of the concept of 'information' in belief revision. Pieces of *information* are inputs to belief states. *Belief states* in turn are the results of a – usually long – history of information processing episodes. In any case, it is plausible to require that the representation of the prior state and the representation of the posterior state should be of the same format.⁹

The format of the input varies in theories of belief revision. The literature has studied pieces of input of the following kinds:

- propositions
- propositions coming with a specification of their relative position in a (total) pre-order
- "ranked" propositions (i.e., propositions coming with ranks)
- ordered pairs of propositions (indicating their comparative retractability)
- propositions with specification of a source
- full preferential structures

Beliefs are usually taken to be represented by propositions – either linguistically, as sentences or sentence-meanings, or abstractly, as sets of possibilities. I shall, somewhat sloppily, not distinguish between these two variants and will identify a sentence, i.e., an expression of a given language, from what is said by such a sentence. Often only the latter is called a 'proposition,' and it is modelled by a set of possible worlds.

A *belief state* determines the set of beliefs held by the agent, but may (and usually does) encode much more information than that. Individual beliefs are in a derivative of belief states. For instance, a belief state may be represented with the help of Grove systems of spheres (Figure 1).

A *system of spheres* is a set of nested sets of possible worlds. The smallest set "in the center" is the set of possible worlds which the agent believes to

⁸Alternatively, one could say that belief revision's "outputs" are the belief states themselves. I shall avoid that terminological move.

⁹This requirement was called the *Principle of categorial matching* in Gärdenfors and Rott (1995).




Figure 1: A Grovean system of spheres

contain the actual world, i.e., the worlds considered “possible” according to the agent’s beliefs. If he receives evidence that the actual world is not contained in this smallest set, he falls back on the next larger set of possible worlds. Thus the first shell¹⁰ around the center contains the worlds considered second most plausible by the reasoner. And again, should it turn out that the actual world is not to be found in this set either, the reasoner is prepared to fall back on her next larger set of possible worlds. And so on. A system of spheres is equivalent to a total pre-ordering of possible worlds (world $w \preceq w'$ if w is contained in every sphere in which w' is contained).

A system of spheres at the same time determines an ordering between propositions, often called *entrenchment ordering*. If A covers each sphere that is covered by B , A is at least as entrenched in the agent’s belief state as B .

The belief state determines the beliefs of the agent by a method of projection or retrieval. If a belief state is represented by a system of spheres (or a total pre-ordering \preceq) of possible worlds, the beliefs are those propositions that are true in each of the possible worlds contained in the innermost sphere (in each of the possible worlds that are minimal under \preceq). If a belief state is represented by an ordering of propositions, say an entrenchment-ordering, the beliefs are those propositions that are non-minimal under this ordering. Belief states in this sense may be regarded as *non-propositional information structures*.

¹⁰A *shell* is the difference set between two neighbouring spheres. Spheres are nested, shells are disjoint. The shells are numbered, but the numbers are not supposed to have any meaning beyond the indication of the ordering of spheres.

2.4 Filling in the parameters

We can now give a first overview of some of the more important stages in the development of belief revision theory. We just need to fill in the parameters into the scheme just specified in various ways (see Table 1).

	input	prior belief state	posterior belief state
<i>AGM 1978ff</i>	proposition	set of beliefs (logically closed; plus preference structure on beliefs or sets of beliefs)	set of beliefs (logically closed)
<i>Grove 1988, Katsuno- Mendelzon 1991</i>	proposition	preference structure on possible worlds	set of possible worlds
<i>Veltman 1976, Kratzer 1981, Nebel 1989, Hansson 1989, 1999</i>	proposition	set of beliefs (syntactically structured, i.e. not logically closed)	set of beliefs (logically closed?)
<i>Spohn 1988, Goldschmidt- Pearl 1992</i>	proposition plus plausibility index	ranking function (a kind of preference structure)	ranking function
<i>Darwiche-Pearl 1994, 1997</i>	proposition (plus plausibility index ?)	general format, with beliefs derivative	general format, with beliefs derivative
<i>Cantwell 1997, Fermé-Rott 2004</i>	pair of sentences	preference structure	preference structure
<i>Nayak 1994ff, merging – fusion</i>	preference structure	preference structure	preference structure

Table 1: Filling in the parameters for a functionalist account of belief change

An important turning point of the development of belief revision theory was the clear recognition in the 1990s that a belief state must not be identified with a (logically closed) set of beliefs. The study of the problem of iterated belief revision made it clear that AGM's belief set had to be replaced by a selection mechanism or a preferential structure in order to encode a full belief

state. Alternatively, Darwiche and Pearl (1994, 1997) suggested that it is best to take “belief state” as a primitive concept. More on this in section 5 below. After the turn of the millenium, much research has focussed on the problem of fusing or merging belief states. The old idea that a belief state (carrying information about the learning history of the agent) gets revised by a single piece of information is no longer valid as a description of the standard problem. The question addressed now is how to merge two or more rather general information structures into a single one.

3 Belief Change = Revision + Reflection

How do belief states change? I suggest to decompose the process of belief change into two different processes. Then there are two fundamentally different approaches, depending on which of the two processes is being highlighted.

The process of *revision* is that of changing the current data base or belief state, in response to the receiving of some new piece of information (some ‘input’). Models of belief change emphasizing the process of revision are taking the *horizontal perspective*.

The process of *reflection* is that of finding an equilibrium state by processing, or drawing inferences from, the currently available information. Models of belief change emphasizing the process of reflection (while having a rather straightforward method os revision) are taking the *vertical perspective*.¹¹

We shall now consider three approaches.

- Foundationalism in vertical perspective
- Foundationalism in horizontal perspective
- Coherentism in horizontal perspective

Coherentism denies that the foundationalists’ distinction between basic and derived beliefs has any clear significance. There is no set of propositions that enjoy the privilege of serving as a foundation for the other beliefs.¹² Coherentists are not interested in the origin of belief states. The reflection component does not even appear in the picture, the aim of reaching (or rather remaining in) an equilibrium state is part and parcel of the revision process (compare Fig. 4 above).

Assume that ‘inputs’ have come in repeatedly. A *data base* is the result of

¹¹The pair ‘reflection’ and ‘revision’ also plays a central role in Harman’s seminal book *Change in View* (1986, Chapter 1), but Harman’s meanings of these terms are different from ours.

¹²Note, however, that nothing in the above foundationalist picture guarantees that the elements of the data base are immune to revision!

collecting the inputs and putting them together. It is important that data bases in our sense need not obey any coherence constraints of a logical or any other nature. A data base will be considered as a rough and ready collection of pieces of information. It is the *basic information structure* figuring in belief revision, directly interpreting or exploiting, as it were, the structure of the inputs.

For instance, data bases can be

- sets of propositions
- totally pre-ordered sets of propositions
- ranked sets of propositions

The difference between totally pre-ordered and ranked sets is that the latter, but not the former, have numbers attached to the propositions signifying their ‘degree of belief.’ These numbers represent distances, and it makes sense to perform arithmetical operations on them.

We assume that in response to incoming input the sets in question may grow or shrink, that there are insertions or deletions at certain positions in the pre-ordered or ranked structure.

The process of *reflection* is that of finding an equilibrium state on the basis of a given data base. The data base is processed and thereby transformed it into a belief state ‘in equilibrium.’ Reflection is ‘static’ in the sense that no ‘new’ input is being dealt with. It can be thought of as an act of *information processing*. Reflection is distinguished from the equally static process of *drawing inferences* which does not yield a belief state, but only a belief set, i.e., a well-balanced set of beliefs (propositions) that can be inferred from, or are supported by, the agent’s data base.¹³ So we distinguish two kinds of “*static*” operations on data bases, and it has turned out that information processing is *more* than the drawing inferences. The latter falls short of the former unless a belief state is identified with a theory to begin with. In any case a belief state determines belief set (=theory)

The process of *revision* is basically a response to incoming input. Revision may consist in a relatively trivial change operation on the data base level, or alternatively, in a relatively sophisticated change operation on the belief state level. Obviously, the nature of a change operation depends on the nature of the entities related by the change operation. On the base level, the changes need not be sophisticated since data bases are not required to be coherent.

If a revision on the level of belief states is induced by a revision on the level of belief bases, reflection and revision have to be in harmony in such a way that

¹³It is plausible to assume two mappings here, one taking data bases to belief states, and another one taking belief states to belief sets. Neither of these mappings is injective. The latter mapping has been called projection or retrieval above.

the diagram in Fig. 2 commutes.




Figure 2: Commuting diagram with reflection and revision

The revision process can be thought of as operating on the *data base level*, thus generating a change operation on the belief state level only derivatively (Fig. 3).




Figure 3: Changes on the data base level

Alternatively, one can think of the revision process as operating on the *belief state level*. Then the lower level of the data bases is not in the picture any more, all deliberation that takes place is part and parcel of the change from one equilibrium state to another one. Information comes as a disturbance of an equilibrium, and processing is just reorganizing belief states in this approach. Since no reflection process is there to eliminate any remaining inconsistency or incoherence, the change process itself has to be sophisticated (Fig. 4). It is important to be clear about the fact that it is not the input alone that brings about the transition from one belief state to another. In order to resolve contradictions or avoid unwanted implications, it is necessary to make choices which beliefs to remove. For non-trivial revision operations, the agent will therefore need a *selection structure* (for instance, a preference relation) to guide these choices, and also a *rule of application* specifying how exactly to apply this structure when accommodating the input.¹⁴

Coherentism may be hidden foundationalism. Often researchers aimed at rep-

¹⁴For instance, the revised belief set by A may be identified with the set of all sentences that are true in all minimal possible worlds that satisfy A , or the set of sentences B for which $A \rightarrow B$ is more entrenched than $\neg A$. But I do not want to say much about concrete rules of application in this paper.




Figure 4: Changes on the belief state level

resentation results of the following sort. Given a revision operation on the belief state level, are there data bases, a reflection operation and a revision operation on the data base level such that the revision operation on belief states is induced by the latter? We will give a few examples indicating that the answer is positive in some interesting cases.

‘Processing information’, or equivalently, ‘processing inputs’ thus may thus mean two different kinds of things. Either the agent *adds* the new piece of information to the current *data base*, to some existing (totally pre-ordered, ranked) set of propositions. Or the agent *accommodates* his or her *belief state* to the input. On the data base level, change operates on the data base level, with a simple and straightforward, unrestrained insertion or deletion of a proposition from a set of propositions. If that set is ordered, addition of a piece of information¹⁵ on top of the ordered list may be suitable, or the new information may come together with some specified ‘rank’ indicating its new position in the existing ranking. Free and simple insertions will usually cause a violation of deductive closure and will sometimes cause a violation of consistency of the information the agent has been provided with. This calls for a sophisticated reflection operation that restores consistency and ultimately produces logical closure. Fig. 5 illustrates the foundationalist idea of the “vertical perspective”.

3.1 Processing data bases: Example 1

Let the *data base* be a set Γ of propositions. The most straightforward method of *drawing inferences* from such a data base is to take the logical closure $Cn(\Gamma)$ of Γ , where the logic Cn used is some broadly classical or Tarskian logic.¹⁶

Using some such standard logic, however, will frequently create *problems of interaction* with simple change operations if the latter are applied to data bases. If a new proposition is simply added to Γ , this may easily generate an in-

¹⁵Or of a ‘phantom belief’, see Rott (2001, p. 129).

¹⁶A Tarskian logic is required to be reflexive, monotonic and to satisfy the deduction theorem.




Figure 5: Foundationalism in the vertical perspective

consistency that cannot be processed by the logic). If a previously accepted proposition is simply eliminated from Γ , this may fail to efficiently remove Γ , since applying logical closure to the remaining propositions may easily generate the cancelled proposition back again. While it is not clear how any remotely standard logic can solve the inconsistency problem, a paraconsistent logic might help here (see Priest, Routley and Norman (1989)).

But one may insist on consistence and effective removal already on the data base level. So many people have argued that it is better to think about change operations on the data base that are not as simple-minded as set-theoretic additions and subtractions from Γ (see Hansson 1999). Instead of just adding or subtracting a single item, additional items have to be removed from the data base in order to avoid lapsing into inconsistency or to avoid that the sentence to be discarded can be derived from the remaining data base. A choice mechanism is necessary for determining which additional items to remove. Fig. 6 gives an impression of how this foundationalist idea which emphasizes the revision process works. It changes from the vertical to the horizontal perspective.

Fig. 7 gives an impression of how the coherentist idea works. By definition, there is no base level in this model, thus no reflection part in the sense we have defined. The idea necessarily emphasizes the revision process, it takes the horizontal perspective.

3.2 Processing data bases: Example 2

Let the *data base* now be a totally pre-ordered set $\langle \Gamma, \prec \rangle$ of propositions, written

$$\Gamma_1 \prec \Gamma_2 \prec \dots \prec \Gamma_n$$




Figure 6: Foundationalism in the horizontal perspective




Figure 7: Coherentism (in horizontal perspective)

where the Γ_i 's are non-empty subsets of Γ .¹⁷

A method of *drawing inferences* from such a data base is described as follows. Construct the maximal consistent subsets X of Γ , maximizing subject to consistency “from the top to the bottom.” That is, first take a maximally consistent subset of Γ_n and keep it; then add in a maximally consistent way elements of Γ_{n-1} and keep this extended set; then add in a maximally consistent way elements of Γ_{n-2} and keep the newly extended set; and so on, until you reach Γ_1 . Call the set obtained by this procedure X . Of course, there are many such X 's, since in general there are many maximally consistent subsets (on each level i). Then there are two inference strategies. If you want to draw *bold* inferences, take *one* such X and close it under your standard logic C_n . Alternatively, if you want to apply *cautious* inferences, take *all* those X 's, close each of them under your standard logic C_n and form the intersection of all the resulting sets. This is what one can infer from the prioritized data base according to the two strategies, i.e., the belief set supported by $\langle \Gamma, \prec \rangle$.

There is no consistency problem generated by the interaction with simple change operations. The ordering of the data base induces as it were some sort of ‘paraconsistent logic.’ But notice that Γ is not in general included in the belief

¹⁷I assume finiteness for simplicity. If the Γ_i 's were allowed to be non-empty, then we would in fact have ordinal numbers in the description of a data base, i.e., we would have a ‘ranked’ data base that allows to express quantitative degrees of belief.

set supported by Γ .¹⁸

There is a way of *processing the database* that corresponds fairly well to this idea of inferences, but it does not give a total pre-ordering as a belief state (cf. Rott 2000).

3.3 Processing data bases: Example 3

Data base: totally pre-ordered set $\langle \Gamma, \prec \rangle$ of proposition

$$\Gamma_1 \prec \Gamma_2 \prec \dots \prec \Gamma_n$$

where the Γ_i 's are non-empty subsets of Γ .

Drawing inferences from $\langle \Gamma, \prec \rangle$: Take $\Gamma_i \cup \dots \cup \Gamma_n$ for i minimal such that $\Gamma_i \cup \dots \cup \Gamma_n$ is consistent.

This is what one can infer from the prioritized data base according to this method, i.e., the belief set supported by $\langle \Gamma, \prec \rangle$. But notice again that Γ is not in general included in this belief set.

Processing the data base: We need to construct a belief state out of the data base $\langle \Gamma, \prec \rangle$, and use both the system of spheres and the entrenchment representation.

Generate a Grovean system of spheres of possible worlds, or more precisely, of models of the language) by taking all the set of models of $\Gamma_i \cup \dots \cup \Gamma_{n-1} \cup \Gamma_n$ for $i = 1, \dots, n$.¹⁹ Or alternatively (and equivalently), generate a total pre-ordering of propositions by defining $A \leq B$ iff for all i , whenever A follows logically from $\Gamma_i \cup \dots \cup \Gamma_n$, so does B .

With these processing mechanisms, there are no consistency and closure problems, no problems of interaction with simple change operations.

4 Representability problems

We return to the idea that coherentism (on the belief state level) may turn out to be (representable as) a form of foundationalism. The following is an application of the method of processing data bases sketched in Section 3.3. The material of this section is put into a much wider perspective in Rott (2006).

¹⁸The closure problem after simple *eliminations* of an element A of the data base can be avoided, if this belief is eliminated by the addition of the “phantom belief” $\neg A$. As a phantom belief, $\neg A$ is counted for the consistency check “from the top to the bottom,” but it is not used in the closure $Cn(X)$. See Rott (2001, Chapter 5).

¹⁹Equivalently, generate a total pre-ordering of possible worlds by defining $w \leq w'$ iff for all i , whenever w' satisfies $\Gamma_i \cup \dots \cup \Gamma_n$, so does w .

4.1 Moderate revision

Let the agent's *belief state* be represented by a Grovean systems of spheres. Suppose that $\neg A$ is accepted in this state, and that the input is A . Then the idea of moderate revision (Nayak 1994, Rott 2003) is as indicated in Figure 8.




Figure 8: Revision by input A

It can be shown that there is a syntactic operation on *belief bases* that corresponds precisely to this semantic operation on systems of spheres. A suitable base

$$\gamma_1 \prec \gamma_2 \prec \dots \prec \gamma_n$$

for the system of spheres gets changed into

$$\gamma_1 \prec \gamma_2 \prec \dots \prec \gamma_n \prec A \prec \gamma_1 \vee A \prec \gamma_2 \vee A \prec \dots \prec \gamma_n \vee A$$

Thus the coherentist ‘moderate’ method of changing belief states turns out to be a hidden form of foundationalism, that is, it is induced by a method of changing a corresponding data base.

4.2 Revision by comparison

Let the agent's belief state again be represented by a Grovean systems of spheres and let \leq be the entrenchment relation generated from this systems as in Section 2.3. Suppose that B is accepted in this state, and that the input is A with the proviso that A should be accepted as firmly as the reference proposition B . Alternatively, the input may be thought of as coming in the form ‘ $B \leq A$.’ Then the idea of revision by comparison (also known as raising, Cantwell 1997 raising) is as indicated in Figure 9.²⁰

²⁰The representation of the idea in terms of changes of entrenchments is rather complicated, see Fermé and Rott (2004).




Figure 9: Revision by A , accepting it at least as strongly as B (input ' $B \leq A$ ')

Again there is a syntactic operation on *belief bases* that corresponds precisely to this semantic operation on systems of spheres. A suitable base

$$\gamma_1 \prec \gamma_2 \prec \dots \prec \gamma_n$$

for the system of spheres gets changed into

$$\gamma_1 \prec \gamma_2 \prec \dots \prec \gamma_{i-1} \prec \gamma_i \wedge A \prec \gamma_{i+1} \prec \gamma_{i+2} \prec \dots \prec \gamma_n$$

where i is chosen such that B follows logically from $\gamma_i \wedge \dots \wedge \gamma_n$, but not from $\gamma_{i+1} \wedge \dots \wedge \gamma_n$.

5 Statics and dynamics of information

Information *as structure* is static, information *as being interpreted* is dynamic. But information is always unitary. So can we conclude that the static picture determines the dynamic one? Another look at the history of belief change theories will help us answer this problem.

From now on, let Φ denote a belief state. The only condition we place on this concept is that it is possible to retrieve the beliefs of an agent from Φ . Let us write $Bel(\Phi)$ for belief set supported by Φ .²¹

Let $\Phi * A$ denote the revised belief state, if Φ is the prior belief state and A is the input (piece of new information).

²¹Usually, we shall think that $Bel(\Phi)$ is closed under some broadly classical logic.

5.1 The 1980s: AGM's classical models

The original models of Alchourrón, Gärdenfors and Makinson (1985) were somewhat ambiguous about the notion of an epistemic state. Officially, an agent's belief state was represented by a *belief set* $\Phi = Bel(\Phi)$, that is, the agent's set of beliefs as represented by sentences of some regimented language. Belief sets were assumed to be closed under the (broadly classical) logic governing this language. As mentioned in Section 3, however, the agent has to make use of a selection structure or preference structure in order to revise his belief set in a reasonable way. *Given such a structure associated with a belief set Φ* , AGM's method (algorithm, recipe, ...) indeed uniquely determines the new beliefs. The resulting belief change function $*$ specifies, for a given belief set Φ , the posterior belief state $\Phi * A$ for any potential input sentence A . It satisfies the following (by now classic) set of *AGM postulates*:²²

(AGM1) $\Phi * A$ is logically closed.

(AGM2) $\Phi * A$ implies A

(AGM3) $\Phi * A$ is a subset of $Cn(\Phi \cup \{A\})$

(AGM4) If A is consistent with Φ , then Φ is a subset of $\Phi * A$

(AGM5) If A is consistent, then $\Phi * A$ is consistent

(AGM6) If $\Phi_1 = \Phi_2$ and A is equivalent with B , then $\Phi_1 * A = \Phi_2 * B$

(AGM7) $\Phi * (A \wedge B)$ is a subset of $Cn((\Phi * A) \cup \{B\})$

(AGM8) If B is consistent with $\Phi * A$, then $\Phi * A$ is a subset of $\Phi * (A \wedge B)$

One problem with the original AGM approach is that it did not provide for *revisions of such selection structures* in response to new information. For this reason, iterated revisions of belief states were largely impossible, and belief revision theory was not fully dynamic.²³

Another problem is that AGM left it open where the selection structures are supposed to come from. They were just assumed to be somewhere in the background, waiting to be explored in belief change processes. In my view, it is hard to imagine an objective measure with which to gauge changes of beliefs. It is much more plausible to assume that the structure guiding the change of an agent's beliefs is part his or her own mental state, and indeed I suggest that it is part of the agent's *belief state*. As such, it will itself be subject to changes in response to new information. This more dynamic way of thinking about the evolution of belief states came to dominate belief revision theory during the 1990s.

²²Given here in a slightly adapted form.

²³A slight adaptation of the AGM definitions, however, allows the same selection structure to be used in the context of arbitrary belief sets. For various ways of implementing this idea, see Alchourrón and Makinson (1985), Areces and Becher (2001) and Rott (2003).

5.2 The 1990s: Iterations

There are two main ways of extending the AGM framework so as to make belief states rich enough to support iterated changes in response to sequences of new information. In both models, change functions do not operate on the agent's *beliefs*, but directly on his or her *belief states*. Arbitrary iterations of belief changes are possible, so a fully developed dynamics becomes feasible.

First, many researchers have suggested to *identify belief states with selection or preference structures*, of the kind that have proven suitable for one-shot AGM belief change. While such a structure is sufficient to uniquely determine the set of current beliefs as well as the AGM revisions of this belief set, it is not sufficient to determine its own revision. A method or rule how to change the selection structure (to apply it to itself as it were) if a new piece of information is comes in has to be specified. Three of the most simple and plausible ideas are surveyed under the names 'radical', 'conservative' and 'moderate' revision in Rott (2003). For instance, the moderate method introduced in Section 4.1 can be fully characterized by adding a single axiom to the AGM postulates that takes care of iterations:²⁴

$$(\text{Mod9}) \quad \Phi * A * B = \begin{cases} \Phi * (A \wedge B) & \text{if } B \text{ is consistent with } A \\ \Phi * B & \text{otherwise.} \end{cases}$$

It looks as if *the statics fully encodes the dynamics of belief*. Each belief state contains all information for all its future revisions. But this is not quite true: In order to perform a change of belief, one needs to specify a certain method, a 'rule of application', like, e.g., the rule for moderate revision.

Let us now turn to the second way of extending the AGM models. In the important paper of Darwiche and Pearl (1994, 1997), a *belief state is introduced as a primitive notion*. A belief state is not identical with a (logically closed) belief set, but the latter is assumed to be retrievable from the belief state with the help of the *Bel* function. The notation of the postulates then needs to be adapted accordingly. Here is the set that these authors proposed:²⁵

- (DP1) $Bel(\Phi * A)$ is logically closed.
- (DP2) $Bel(\Phi * A)$ implies A
- (DP3) $Bel(\Phi * A)$ is a subset of $Cn(Bel(\Phi) \cup \{A\})$
- (DP4) If A is consistent with $Bel(\Phi)$, then $Bel(\Phi)$ is a subset of $Bel(\Phi * A)$
- (DP5) If A is consistent, then $Bel(\Phi * A)$ is consistent
- (DP6) If $\Phi_1 = \Phi_2$ and A is equivalent with B , then $Bel(\Phi_1 * A) = Bel(\Phi_2 * B)$
- (DP7) $Bel(\Phi * (A \wedge B))$ is a subset of $Cn(Bel(\Phi * A) \cup \{B\})$

²⁴Also compare Nayak (1994) and Nayak, Pagnucco and Peppas (2003).

²⁵Also in a slightly adapted form, in order to facilitate the comparison with AGM.

(DP8) If B is consistent with $Bel(\Phi * A)$, then $Bel(\Phi * A)$ is a subset of $Bel(\Phi * (A \wedge B))$

The only difference with standard AGM belief revision postulates resides in Darwiche-Pearl's sixth condition:²⁶

(AGM6') If $Bel(\Phi_1) = Bel(\Phi_2)$ and A is equivalent with B , then $Bel(\Phi_1 * A) = Bel(\Phi_2 * B)$

Darwiche and Pearl (1997) add four postulates for the iterated revision of belief states.

(DP9) If B implies A , then $Bel((\Phi * A) * B) = Bel(\Phi * B)$

(DP10) If A is inconsistent with B , then $Bel((\Phi * A) * B) = Bel(\Phi * B)$

(DP11) If A is implied by $Bel(\Phi * B)$, then it is implied by $Bel((\Phi * A) * B)$

(DP12) If A is consistent with $Bel(\Phi * B)$, then it is consistent with $Bel((\Phi * A) * B)$

These postulates have a convincing semantic motivation. But like the AGM postulates, they do not specify a unique method for changes of belief states. Only *given* a preference structure and *given* a specific rule of application (like, e.g., that of moderate revision), is the posterior belief state determinately fixed.

We said above that it is plausible to regard preference structures as parts of the agent's belief state. But where do rules of application come from? Preferences are non-propositional, but still they seem to represent something. Rules of application, on the other hand, are not 'declarative' in any sense. They constitute 'procedural' information about how to apply preference structures in the process of rebuilding one's belief state. Could rules of application perhaps belong to the agent's belief state, too? These questions have not been addressed in the belief revision literature yet, but they do present interesting and important challenges from a philosophical point of view.

5.3 The 2000s: Merging and fusion

From the end of the past century on, research on belief change has become more and more focused on the fusion (also known as the merging, combination, integration, arbitration, ...) of belief states.²⁷

²⁶It is the fourth condition in Darwiche and Pearl's numbering, see Darwiche-Pearl (1997), p. 7.

²⁷Among the many relevant papers are Baral, Kraus, Minker and Subrahmanian (1992), Revesz (1993), Nayak (1994), Konieczny and Pino-Perez (1998), Liberatore and Scherf (1998), Benferhat, Dubois, Prade and Williams (1999), Meyer (2000), Andreka, Ryan and Schobbens (2002) and Liao (2005).

In traditional theories of belief change, the input was usually treated as a single piece of information. In belief fusion, the ‘input’ is one or more data bases or belief states of other agents. Earlier there had been a clear asymmetry between input and data base. The former was called ‘new information’, the latter is some form of representation of the result of the previous information that an agent had received and processed.²⁸ In belief fusion, no such asymmetry is assumed, although it may of course be stipulated as a special constraint on a fusion problem. Belief revision in the older style may thus be viewed as a special case of belief fusion. With the turn to belief fusion, the area of belief revision has left the restrictions of a single agent environment and moved into a multi-agent setting. Now multiple belief states can be dealt with.

This field is extremely active today, and it would be presumptuous here to try to survey the diversity of paths followed in belief fusion. I rather try to convey the flavour of any such undertaking by presenting the axiomatic characterization of Konieczny and Pino-Perez’s (1998) account. By a *belief base*, these authors mean just a single proposition A representing (the conjunction of) a person’s beliefs. By a *belief set*, they denote a multiset $E = [A_1, \dots, A_n]$ of propositions, where A_i is the belief set of the i th person. $\Delta(E)$ denotes the belief base that results from merging the elements of E . Here are the postulates for *merging* or *fusion* suggested by Konieczny and Pino-Perez:

- (KP1) $\Delta(E)$ is a consistent proposition
- (KP2) If the belief sets A_1, \dots, A_n in E are jointly consistent, then $\Delta(E) = A_1 \wedge \dots \wedge A_n$
- (KP3) If E_1 is elementwise equivalent with E_2 , then $\Delta(E_1) = \Delta(E_2)$
- (KP4) If A_1 is inconsistent with A_2 , then $\Delta([A_1] \sqcup [A_2]) \not\vdash A_1$
- (KP5) $\Delta(E_1) \wedge \Delta(E_2)$ implies $\Delta(E_1 \sqcup E_2)$
- (KP6) If $\Delta(E_1)$ is consistent with $\Delta(E_2)$, then $\Delta(E_1 \sqcup E_2)$ implies $\Delta(E_1) \wedge \Delta(E_2)$

For motivation and semantics, we have to refer the reader to the original paper.

Given the idea that belief states are to be identified with preference structures suitable for resolving the potential conflicts between different units of information, it is not surprising that a large part of the problems involved in belief fusion present themselves as problems of amalgamating or aggregating preference relations. This is a general problem that can be considered in abstraction from the problems specifically pertaining to information processing.

Natural links are established with social choice theory, game theory, negotiation theory, etc.

²⁸Perhaps together with whatever constitutes the agent’s a priori beliefs.

6 Conclusion

We have seen different kinds of information structures at work in belief revision. There is propositional information (beliefs and inputs), there is non-propositional information that still, in some sense, seems to represent something (the preference orderings that we have identified with belief states), and there is non-propositional, non-representing, purely procedural information (rules of application specifying how to use the preference orderings in the process of belief revision).

We have further seen different models of belief change. In the classical models, beliefs were determined by preferences and rules of application for the use of these preferences. In the 1990s, preferences themselves were determined by prior preferences and rules of application for the change of these preferences. The question concerning the choice or change of rules of application, however, has remained unanswered. Ultimately, this brings us to the question whether believers are free to ‘use’ information as we like. Do we, in this sense, possess ‘informational freedom’? Or is, in the picture afforded by the literature on belief revision, everything about our beliefs determined?

The philosopher Galen Strawson (1986, 1994) has put forward a well-known argument to the effect that there can be no free, responsible action which goes as follows. Any free action of a person P (that is, any action for which she is responsible, any action performed for a reason) is a consequence, among other things, of ‘the way P is, mentally speaking’ or, as Strawson also puts it, her ‘mental nature’ or ‘character.’ Thus person P is responsible for her action only if she is responsible for her character. She can intentionally choose her character only if she is equipped with ‘principles of choice, $P1$ - with preferences, values, pro-attitudes, ideals – in the light of which [she] chooses how to be’ (1994, p. 7). Person P is responsible for her character only if she is responsible for her principles of choice $P1$. She is responsible for the latter only if she has intentionally chosen them, which in turn is possible only if she is equipped with second-order principles of choice $P2$ for the choice of her first-order principles of choice $P1$. And so on, *ad infinitum*. Strawson concludes: ‘True self-determination is impossible because it requires the actual completion of an infinite series of choices of principles of choice.’ (1994, p. 7) Therefore, there can be neither true freedom nor true responsibility.

What is interesting about this argument from our point of view is that Strawson’s a priori argument appears to describe quite exactly what has in fact happened in the development of belief revision theory. The agent’s changes of beliefs are determinate, provided that higher-order information structures (preferences and rules of application) of the belief state are given. What helps to solve the belief change problem at one level, however, is itself subject to revision. The changes of the preference-orderings (often used as models of belief states) are themselves determinate, provided that some method of preference change is assumed as given. But so far no choice mechanisms for the choice of

rules of application have been proposed. And provided that we had such choice mechanisms, how would *they* be rationally selected? Should we assume that there is freedom of choice here, or is this just a matter of having such-and-such a mental nature or character? These are not questions that haunt computer scientists and information technicians in their daily work, but they are hard to dismiss from a philosophical point of view.

References

Carlos Alchourrón and David Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.

Carlos Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50:510–530, 1985.

Hajnal Andréka, Mark Ryan, and Pierre-Yves Schobbens. Operators and laws for combining preference relations. *Journal of Logic and Computation*, 12:13–53, 2002.

Carlos Areces and Verónica Becher. Iterable AGM functions. In Mary-Anne Williams and Hans Rott, editors, *Frontiers in Belief Revision*, pages 261–277. Kluwer, 2001.

Alexandru Baltag, Hans van Ditmarsch, and Larry Moss. Epistemic logic and information update. In Pieter Adriaans and Johan van Benthem, editors, *Handbook of the Philosophy of Information*. Elsevier, 2006.

Chitta Baral, Sarit Kraus, Jack Minker, and V.S. Subrahmanian. Combining multiple knowledge bases consisting of first order theories. *Computational Intelligence*, 8:45–71, 1992.

Salem Benferhat, Didier Dubois, Henri Prade, and Mary-Anne Williams. A practical approach to fusing prioritized knowledge bases. In Pedro Barahona and Jos Jlio Alferes, editors, *Progress in Artificial Intelligence. 9th Portuguese Conference on Artificial Intelligence, EPIA '99, vora, Portugal, September 21-24, 1999*, volume 1695 of *LNAI*, pages 222–236. Springer, 1999.

John Cantwell. On the logic of small changes in hypertheories. *Theoria*, 63:54–89, 1997.

Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. In R. Fagin, editor, *TARK'94 – Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 5–23, San Mateo, CA, 1994. Morgan Kaufmann.

Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.

Eduardo Fermé and Hans Rott. Revision by comparison. *Artificial Intelligence*, 157:5–47, 2004.

André Fuhrmann. On the modal logic of theory change. In André Fuhrmann and Michael Morreau, editors, *The Logic of Theory Change*, LNCS 465, pages 259–281. Springer, Berlin, 1991.

Peter Gärdenfors and Hans Rott. Belief revision. In Dov M. Gabbay, Christopher J. Hogger, and John A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume Volume IV: Epistemic and Temporal Reasoning, pages 35–132. Oxford University Press, Oxford, 1995.

Peter Gärdenfors. Conditionals and changes of belief. In Illkka Niiniluoto and Raimo Tuomela, editors, *The Logic and Epistemology of Scientific Change*, volume 30, nos. 2–4 of *Acta Philosophica Fennica*, pages 381–404. North Holland, Amsterdam, 1978.

Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge, Mass., 1988.

Moises Goldszmidt and Judea Pearl. Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, pages 661–672, Cambridge, MA, 1992. Morgan Kaufmann.

Jeroen Groenendijk and Martin Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–101, 1991.

Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.

Sven Ove Hansson. New operators for theory change. *Theoria*, 55:114–132, 1989.

Sven Ove Hansson. *A Textbook on Belief Dynamics*. Kluwer, Dordrecht, 1999. With a booklet "Solutions to Exercises", Kluwer Academic Publishers, Dordrecht 1999.

Gilbert Harman. *Change in View*. Bradford Books, MIT Press, Cambridge, Mass, 1986.

Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

Kevin Kelly, Oliver Schulte, and Vincent Hendricks. Reliable belief revision. In Maria Luisa Dalla Chiara, Kees Doets, Daniele Mundici, and Johan van Benthem, editors, *Logic and Scientific Methods - Proceedings of the 10th International Congress of Logic, Methodology and Philosophy of Science*, pages 383–398. Kluwer, Dordrecht, 1997.

Kevin Kelly. Iterated belief revision, reliability, and inductive amnesia. *Erkenntnis*, 50:11–58, 1999.

Jaegwon Kim. *Philosophy of Mind*. Westview Press, Boulder, Colorado, second edition, 2006.

Sébastien Konieczny and Ramón Pino-Pérez. On the logic of merging. In A. G. Cohn, L. Schubert, and S. C. Shapiro, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixth International Conference (KR'98)*, pages 488–498, San Francisco, CA, 1998. Morgan Kaufmann.

Angelika Kratzer. Partition and revision: The semantics of counterfactuals. *Journal of Philosophical Logic*, 10:201–216, 1981.

Churn-Jung Liau. A modal logic framework for multi-agent belief fusion. *ACM Transactions on Computational Logic (TOCL)*, 6:124–174, 2005.

Paolo Liberatore and Marco Scherf. Arbitration (or how to merge knowledge bases). *IEEE Transactions Knowledge and Engineering*, 10:76–90, 1998.

Sten Lindström and Włodzimierz Rabinowicz. Ddl unlimited: Dynamic doxastic logic for introspective agents. *Erkenntnis*, 50:353–385, 1999.

David Makinson. Ways of doing logic: What was different about AGM 1985? *Journal of Logic and Computation*, 13:3–13, 2003.

Thomas Meyer. Merging epistemic states. In Riichiro Mizoguchi and John Slaney, editors, *PRICAI 2000: Topics in Artificial Intelligence*, volume 1886 of *LNAI*, pages 286–296. Springer-Verlag, Berlin, 2000.

Abhaya Nayak, Maurice Pagnucco, and Pavlos Peppas. Dynamic belief revision operators. *Artificial Intelligence*, 146:193–228, 2003.

Abhaya Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.

Bernhard Nebel. A knowledge level analysis of belief revision. In Ronald Brachman, Hector Levesque, and Ray Reiter, editors, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 301–311, San Mateo, CA, 1989. Morgan Kaufmann.

Grahm Priest, Richard Routley, and Jean Norman (eds.). *Paraconsistent Logic. Essays on the Inconsistent*. Philosophia Verlag Mnchen, München, 1989.

Hilary Putnam. Minds and machines. In Sidney Hook, editor, *Dimensions of Mind*, pages 138–164. State University of New York Press, Albany, N.Y., 1960. Reprinted in Hilary Putnam, *Mind, Language and Reality*, Cambridge 1975, pp. 362–385.

Hilary Putnam. Psychological predicates. In W.H. Captain and D.D. Merrill, editors, *Art, Mind and Religion*, pages 37–48. University of Pittsburgh Press, 1967. Reprinted as ‘The Nature of Mental States’, in Hilary Putnam, *Mind, Language and Reality*, Cambridge, 1975, pp. 429–440.

Peter Z. Revesz. On the semantics of theory change: Arbitration between old and new information. In *Proceedings PODS'93: 12th ACM SIGACT SIGMOD SIGART Symposium on the Principles of Database Systems*, pages 71–82, 1993.

Hans Rott. “Just because”: Taking belief bases seriously. In Samuel R. Buss, Petr Hájek, and Pavel Pudlák, editors, *Logic Colloquium '98 – Proceedings of the Annual European Summer Meeting of the Association for Symbolic Logic held in Prague*, volume 13 of *Lecture Notes in Logic*, pages 387–408. Association for Symbolic Logic, Urbana, Ill., 2000.

Hans Rott. *Change, Choice and Inference: A Study in Belief Revision and Nonmonotonic Reasoning*. Oxford University Press, Oxford, 2001.

Hans Rott. Coherence and conservatism in the dynamics of belief. Part II: Iterated belief change without dispositional coherence. *Journal of Logic and Computation*, 13:111–145, 2003.

Hans Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In Henrik Lagerlund, Sten Lindström, and Rysiek Sliwinski, editors, *Modality Matters: Twenty-Five Essays in Honour of Krister Segerberg*, pages 359–384. Uppsala Philosophical Studies, Vol. 53, Uppsala Universitet, 2006.

Krister Segerberg. The basic dynamic doxastic logic of AGM. In Mary-Anne Williams and Hans Rott, editors, *Frontiers in Belief Revision*, pages 57–84. Kluwer, Dordrecht, 2001.

Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, Princeton, New Jersey, 1976.

Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W.L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, pages 105–134. Kluwer, Dordrecht, 1988.

Robert C. Stalnaker. *Inquiry*. Bradford Books, MIT Press, Cambridge, MA, 1984.

Galen Strawson. *Freedom and Belief*. Clarendon Press, Oxford, 1986. Reprinted with corrections 1991.

Galen Strawson. The impossibility of moral responsibility. *Philosophical Studies*, 75:5–24, 1994. (Reprinted in Gary Watson, ed., *Free Will*, Oxford University Press, 2nd edition 2002, pp. 212–228).

Alan Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.

Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, to appear. ILLC Research Reports, PP-2006-11, University of Amsterdam, 2006.

Frank Veltman. Prejudices, presuppositions and the theory of counterfactuals. In Jeroen Groenendijk and Martin Stokhof, editors, *Amsterdam Papers of Formal Grammar*, volume 1, pages 248–281, 1976.

Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25:221–261, 1996.