

# Two Dogmas of Belief Revision

Hans Rott

Institute for Philosophy, University of Regensburg,  
93040 Regensburg, Germany  
hans.rott@psk.uni-regensburg.de

## Abstract

The paper attacks the widely held view that belief revision theories, as they have been studied in the past two decades, are founded on the Principle of Informational Economy. The principle comes in two versions. According to the first, an agent should, when accepting a new piece of information, aim at a minimal change of his previously held beliefs. If there are different ways to effect the belief change, then the agent should, according to the second version, give up those beliefs that are least entrenched. Although widely proclaimed by belief revision theorists, I argue that both versions of the principle are dogmas that are not (and perhaps should not be) adhered to. I substantiate this claim with two simple mathematical observations, and I defend it against four possible objections that involve contractions, reconstructions, dispositions, and truths.

## 1. Introduction

Quine's paper "Two Dogmas of Empiricism" (1951) was epoch-making in philosophy. It sealed the fate of logical positivism by undermining two presumptions of classical as well as modern empiricism: The idea that there is a sharp distinction between analytical judgements (knowledge of meanings) and synthetical judgements (knowledge of facts), and the idea that every meaningful sentence can be reduced to a construction upon observation reports. The interesting point for us is that Quine closes his paper with a beautiful section on "empiricism without the dogmas" which deals almost exclusively with a topic that would nowadays be called *theory change*, or *belief revision*. He paints a picture of how we should, and for the most part do, accommodate our scientific heritage if we meet with recalcitrant experiences. According to Quine's pragmatist approach, belief revision is a matter of choice, and the choices are to be made in a such a way that (a) the resulting belief system squares with the experience, (b) it is simple and (c) the choices disturb the original system as little as possible.

Quine's picture is mainly metaphorical, and he never intended to turn his picture into a formal theory of theory change. The situation today is different. At least since the seminal work of Alchourrón, Gärdenfors and Makinson (AGM for short) in the early 1980s, there has been a clear-cut paradigm of how to logically formalize theory change (Alchourrón, Gärdenfors and Makinson 1985, Gärdenfors 1988, Gärdenfors and Rott 1995, Hansson 1999a). The original AGM approach has turned out to be limited in various ways, and many extensions and revisions of it have been developed over the past two decades. In this paper, I do not want to find fault with any of the formal theories of belief revision that are currently being advocated. My aim is rather to call in question what has been taken to be *the* principal idea behind the current theories of belief revision: The idea of *informational economy*. This idea is basically the same as Quine's criterion (c), often also called the *principle of minimum mutilation* or *conservatism*. I call it in question not in itself, but as *the* philosophical background that makes AGM style theories of belief revision intelligible.

The principle on informational economy tells us that we should not give up our beliefs beyond necessity. Now it is not spelt out clearly what "necessity" here means. It might be suggested that we can satisfy the principle of minimal change (or inertia) in an ideal way by not changing our beliefs at all?<sup>1</sup> But that is certainly not intended. There are basic requirements for belief revision that have to be satisfied, and only against the background of these conditions can we apply the minimal change principle. One of these basic requirements that we are going to endorse in this paper is that revisions should be *successful* in accommodating new information: The kind of revision that we consider is such that the incoming input actually gets accepted.<sup>2</sup>

The second basic requirement we are going to impose is that agents should be taken to be ideally competent as regards logic. They should accept all the consequences of the beliefs they hold (that is, their set of beliefs should be *logically closed*), and they should see to it that their beliefs are *consistent* (and it is just the task of belief revision theory to give an account of *how* they should see to it that their beliefs remain consistent).<sup>3</sup> I shall assume that the *basic requirements* – that belief sets should be closed and that revisions should be successful and lead to consistent belief sets – form the background against which questions of informational economy or minimal change are discussed. Agents who do not change their beliefs at all can perhaps be called very economical in the administration of what they currently possess, but they would not be able to

---

<sup>1</sup>Thanks to Isaac Levi for raising this point.

<sup>2</sup>Of course this is too strict a constraint for a truly general and realistic theory of belief change. This has been recognized by many writers and is taken account of in recent papers on "non-prioritized belief revision." See in particular the special issue of the journal *Theoria* edited by Sven Ove Hansson (1999b).

<sup>3</sup>Again, as stressed by researchers in belief *base* revision and paraconsistent logics, the insistence on closure and consistency may be regarded as unrealistic in general.

successfully interact with the world. “Do nothing!” is not a feasible option for our cognitive life.

## 2. The dogmas

The idea of informational economy has been proclaimed to lie at the basis of belief change from the very beginning of the systematic study of belief revision. The dogma comes in two distinct versions:

- (1) When accepting a new piece of information, an agent should aim at a minimal change of his old beliefs.
- (2) If there are different ways to effect a belief change, the agent should give up those beliefs that are least entrenched.

These maxims have been accompanying belief change theory since its inception, and they have been repeated in the literature time and again. Here is some evidence for this claim.

Ad (1).

“The concept of contraction leads us to the concept of *minimal change of belief*, or briefly, *revision*.” (Makinson 1985, p. 352)

“The criterion of informational economy demands that as few beliefs as possible be given up so that the change is in some sense a *minimal* change of  $K$  to accommodate for  $A$ .” (Gärdenfors 1988, p. 53)

“The amount of information lost in a belief change should be kept minimal.” (Gärdenfors and Rott 1995, p. 38)

“At the centre of the AGM theory [of theory change] are a number of approaches to giving formal substance to the maxim [of minimal mutilation: Keep incisions into theories as small as possible!].” (Fuhrmann 1997, p. 17)

“The hallmark of the AGM postulates is the principle of minimal belief change, that is, the need to preserve as much of earlier beliefs as possible and to add only those beliefs that are absolutely compelled by the revision specified.” (Darwiche and Pearl 1997, p. 2)

Ad (2).

“When a belief set  $K$  is contracted (or revised), the sentences in  $K$  that are given up are those with the *lowest* epistemic entrenchment.” (Gärdenfors 1988, p. 87)

“The guiding idea for the construction is that when a knowledge system  $K$  is revised or contracted, the sentences in  $K$  that are given up are those having the *lowest* degrees of epistemic entrenchment.” (Gärdenfors and Makinson 1988, p. 88)

“In so far as some beliefs are considered more important or entrenched than others, one should retract the least important ones.” (Gärdenfors and Rott 1995, p. 38)<sup>4</sup>

“... when it comes to choosing between candidates for removal, the least entrenched ones ought to be given up.” (Fuhrmann 1997, p. 24)

“A hallmark of the AGM theory is its commitment to the principle of *informational economy*: beliefs are only given up when there are no less entrenched candidates. ... If one of two beliefs must be retracted in order to accommodate some new fact, the less entrenched belief will be relinquished, while the more entrenched persists.” (Boutilier 1996, p. 264–265)

Although (1) and (2) look quite different and it is not immediately clear how they fit together, there is a result that appears to show that they are even “at root identical”<sup>5</sup> and that they can therefore be viewed as two incarnations of a unified idea of informational economy. The result mentioned maps onto one another two important types of belief change constructions investigated by Alchourrón, Gärdenfors and Makinson: belief changes obtained by partial meets of maximal non-implying sets of beliefs (Alchourrón, Gärdenfors and Makinson 1985), and belief changes based on epistemic entrenchment (Gärdenfors and Makinson 1988). These two methods are generally taken to be closely associated with (1) and (2), respectively, and so it is both surprising and pleasing to find that they can be proved equivalent in a rather strict sense, by directly relating the underlying relations used.<sup>6</sup>

The overall picture that we have now gotten of the AGM approach seems nice and harmonious. But I want to argue in this paper that *maxims (1) and (2) are a travesty of the principles that have really been followed in the traditional theories of belief revision*. The philosophical underpinnings of the prevailing theories are not what the folklore would like to have them.

As will become clear later on, I call (1) and (2) “dogmas” *not* because almost all researchers *kept* to these principles (quite the opposite is true), but because so many authoritative voices have *proclaimed* these principles to be the philosophical or methodological rationale for their theories.

---

<sup>4</sup>Similarly Gärdenfors (1992, pp. 9, 17).

<sup>5</sup>As, by the way, Quine (1951, Section 5) claimed the “two dogmas of empiricism” were.

<sup>6</sup>The result is presented in Rott (1991a, Section 4). I shall return to this topic in subsection 4.2 and argue that the interpretation just given is wrong. Criteria (1) and (2) may even be pulling in opposite directions, compare footnote 18.

### 3. Against the dogmas

The following simple fact will be used to attack principle (1). It only draws on the basic requirements for belief revision. Irrespective of the prior belief set  $K$ , I shall call any consistent and logically closed set including a sentence  $\phi$  a *candidate revision* of  $K$  by  $\phi$ .

*Observation 1.* No two distinct belief-contravening candidate revisions of a consistent and logically closed belief set by a sentence  $\phi$  can be set-theoretically compared in terms of the sets of beliefs on which they differ with the prior belief set.<sup>7</sup>

(Proofs of all observations but one are given in the appendix to this paper.) Observation 1 shows that the most straightforward idea to measure conservativity in terms of a comparison of the sets of beliefs in which original and revised theories differ fails.

In this sense, *all* (closed and successful) candidate revisions of a consistent belief set are minimally removed from it. Maxim (1) therefore cannot be a good recommendation. In one reading, it can be used to licence the choice of any arbitrary (successful and closed) revision. In another reading, one which recommends the cautious strategy of taking the greatest common denominator of all minimally deviant sets, it has as a result that all belief-contravening belief changes are amnesic.<sup>8</sup> Here I call a revision *amnesic* if the revised belief set consists of nothing else but the logical consequences of the sentence to be revised with; otherwise we call it *anamnesic*.

This seems to be a strong indication that in order to make good sense of the idea of informational economy we need to turn to a refined description of belief states, such as the one afforded by an ordering of beliefs in terms of their epistemic entrenchment. The intuition behind the term ‘entrenchment’, in the sense that has been given a formal analysis in the literature on belief revision, is basically that of a relation of comparative retractability. A belief  $\phi$  is *more entrenched* than another belief  $\psi$  if and only if the agent holds on to  $\phi$  and gives up  $\psi$  after learning (or hypothetically assuming) that not both  $\phi$  and  $\psi$  are (can be, may be) true.

Having a rudimentary understanding of the term ‘entrenchment’ as it is used in the AGM tradition, we now turn to the principle (2). Like the previous observation, the following one is mathematically trivial, but conceptually striking. We call a new piece of information *moderately surprising* if its negation is an element of the original belief set which is not minimally entrenched in that set.

---

<sup>7</sup>This observation bears some similarity with Miller’s (1974) and Tichý’s (1974) celebrated refutations of Popper’s early concept of verisimilitude. While Popper, Miller and Tichý are interested in the distances of (false) theories from the theory containing all and only the truths, we are interested in the distances of (belief-contraveningly revised) candidate theories from a given prior theory.

<sup>8</sup>This result was proven as Observation 2.2 in Alchourrón and Makinson (1982).

*Observation 2.* Suppose we want to revise a belief set by a sentence  $\phi$  and find two elements of the belief set that non-redundantly entail the negation of  $\phi$ . Then it may well be rational, according to the standard belief revision constructions, to restore consistency by removing the more entrenched and retain the less entrenched belief. In fact, such a situation can always be identified in an anamnestic revision by a consistent and moderately surprising sentence.<sup>9</sup>

Principle (2) thus turns out to be plainly wrong on the “local” interpretation (entailment sets with only two elements) applied in this observation. Actually, the proof of Observation 2 does not presume that belief revision is effected according to the traditional Gärdenfors-Makinson (1988) construction recipe for contractions, but it does draw on the logical structuring of the entrenchment relations they introduce. Figure 1 gives a simple illustration of the situation described in Observation 2, using the Grovean representation of AGM-style belief change.<sup>10</sup>




Figure 1: Example for Obs. 2

The situation is this: two beliefs,  $q$  and  $q \supset \neg p$ , form a set minimally implying the

<sup>9</sup>What is excluded by the definition of entrenchment advocated in Gärdenfors and Makinson (1988) and generalized in Rott (1992b), is that the situation of Observation 2 arises when the two “culprits”  $\psi$  and  $\chi$ , say, *exactly* entail  $\neg\phi$ , that is, when  $\psi \wedge \chi$  is logically equivalent with, but not logically stronger than  $\neg\phi$ . For then we have, by definition,  $\psi < \chi$  if and only if  $\chi$  is, but  $\psi$  is not, in the contraction of the belief set with respect to  $\neg\phi$ , and therefore  $\chi$  is, but  $\psi$  is not, in the revision of the original belief set with respect to  $\phi$ .

<sup>10</sup>I am using the possible worlds representation of Grove (1988) (compare Gärdenfors 1988, pp. 83–86, 94–97). In the center we find the worlds satisfying all current beliefs. The entrenchment of a belief  $\phi$  is measured by the distance of the closest  $\neg\phi$ -worlds from the center. The result of revising the current belief set by  $\psi$  is the set of all sentences satisfied by all closest  $\psi$ -worlds.

negation of the incoming information  $p$ ;  $q$  is strictly less entrenched than  $q \supset \neg p$ ; and yet the agent keeps the less entrenched  $q$  and abandons the more entrenched  $q \supset \neg p$  – and all this according to the accepted AGM methodology.<sup>11</sup>

The situation sketched here is explained in more detail in the appendix. I believe that the paradoxical appearance of Observation 2 is not due to a shortcoming of the AGM recipe, but due to the fact that principle (2) expresses a wrong ideology (or, perhaps, the fact that it gives a wrong expression to an unclear idea<sup>12</sup>).

#### 4. Four attempts to save the dogmas, and why they fail

In the previous section I have been attacking the folklore idea that informational economy or minimum mutilation is the basis of what have long been taken to be the most promising approaches to belief change. The reason for this was that the idea of minimal change is difficult to formulate (Observation 1), or that its application in the construction of revisions is ill-understood (Observation 2).

I am now going to briefly discuss four attempts to defend the minimal change idea against the above criticism. The first defence says that I have misdirected my attack by addressing it to belief revisions rather than to the more fundamental operations of belief contraction for which even two ways of characterizing minimality seem possible. The second defence counters my arguments by saying that I have been forgetful of nothing less than the core of the AGM theory, to wit, the representation theorems linking the well-known AGM rationality postulates to the existence of orderings that can be used for determining a reasonable standard of minimality. The third defence advances the claim that if AGM have not been successful in circumscribing minimality, this was so only because their notion of a doxastic state was too simplistic, and as soon as we introduce richer models that in addition to plain beliefs also represent dispositions for belief change, we find a natural way to make conservatism both comprehensible and defensible. The fourth defence argues that we should take up the connection of beliefs with the real world that has been lost by the usual “internalistic” approaches to belief

---

<sup>11</sup>Here is an objection. Observation 2 uses an entrenchment relation not in the GM way, but in the way in which Alchourrón and Makinson (1985) use hierarchies in so-called safe contractions. And it’s no wonder that such a misapplication should lead to unexpected results. – Rejoinder: A result of Rott (1992a, Theorem 4(ii)) shows that entrenchment relations can be used for constructing safe contractions and lead to exactly the same results as when applied in the “proper” GM way. But, how can this be? Wouldn’t a safe contraction clearly eliminate  $q$  and keep  $q \supset \neg p$  in the example used for the proof of Observation 2? No. The belief  $q \supset \neg p$  is lost as the minimal element in another entailment set (viz.,  $\{p \supset q, q \supset \neg p\}$ ), and  $q$  – while eliminated as the minimal element of the set  $\{q, q \supset \neg p\}$  – is ultimately rederived from the “safe” elements  $p \supset q$  and  $p \vee q$ .

<sup>12</sup>One could try to find a more “global” formulation of the idea.

change. If we respect truth as an aim of inquiry in our characterizations of minimal change, the disturbing results may be expected to disappear..

I shall argue that all these attempts to save the dogmas, though promising at first sight, are mistaken.

#### 4.1. Contractions

One attempt to save the idea of informational economy as a cornerstone of the traditional theories of belief change is to recall that most of the time AGM and their followers have written about belief *contractions*, i.e., the problem of rationally giving up a belief. Any contracted belief set is by definition a subset of the prior set, so here the problem with symmetric differences (that made it possible to prove Observation 1) does not arise. Thus we have an obvious standard by which to measure differences. If a candidate contraction<sup>13</sup>  $K'$  is a proper subset of another candidate contraction  $K''$ , then  $K'$  is farther removed from the original set  $K$  than  $K''$ , since the set-theoretic difference  $K - K'$  properly includes  $K - K''$ . An early suggestion of Isaac Levi's is that each legitimate revision by  $\phi$  must be decomposable into a contraction of  $K$  by  $\neg\phi$ , followed by an addition of  $\phi$  and logical closing-up (this is the so-called Levi identity). In this way distances of revisions from the prior set would be expressible in terms of well-defined distances of the corresponding contractions from the prior set. However, this is not intuitively plausible. If  $K'$  and  $K''$  are two candidate contractions with respect to  $\neg\phi$  and if  $K'$  is a subset of  $K''$ , then the result of adding  $\phi$  to both of them and closing them up under logical consequence leads to a belief set  $Cn(K' \cup \{\phi\})$  that is a subset of the belief set  $Cn(K'' \cup \{\phi\})$ . But is the latter closer to  $K$  than the former? There seems to be no justification for saying so. While the latter set contains original beliefs that may have been lost in the former, it also contains novel beliefs that were not held before and may therefore be considered to be gratuitous additions. This fact is basically what Observation 1 is about.

Another advantage of focussing on contractions might be seen in the fact that AGM work with the so-called postulate of Recovery for contractions. That postulate says that when contracting with respect to some sentence  $\psi$ , we may only withdraw beliefs to such an extent that adding back  $\psi$  immediately after the contraction will make us recover all the original beliefs. This postulate may indeed be regarded as an explicit condition of minimum mutilation. However, as a general defence of minimum mutilation, the argument fails for three reasons. First, Recovery is at best a partial encoding of informational economy, since it does not even disallow amnesic belief change.<sup>14</sup> Second,

---

<sup>13</sup>A *candidate contraction* of a belief set  $K$  with respect to a sentence  $\phi$  is a subset of  $K$  which is logically closed and does not contain  $\phi$ .

<sup>14</sup>In terms of the Grovean possible worlds model, recovery disallows the gratuitous admission of  $\phi$ -



the information-preserving effects of Recovery evaporate completely when contractions are used as intermediate steps in the construction of belief revisions with the help of the Levi identity. Third, in many situations meeting the requirement of Recovery has counterintuitive consequences. This postulate has accordingly suffered from severe criticism from numerous members of the belief revision community and cannot be considered as belonging to the core of traditional belief revision theory any longer.

We conclude that neither of the two arguments to support the idea of minimal change through a consideration of belief contractions is convincing.

## 4.2. Reconstructions

The success of the program of Alchourrón, Gärdenfors and Makinson was not in the first place based on their putting together a list of “rationality postulates”, but on their showing that the belief change behaviour thus axiomatized can be represented by a number of interesting and plausible explicit constructions (such as the partial meet constructions and the entrenchment-based contractions mentioned above). All of these constructions make use of some kind of structure of the belief state that guides the selection of most plausible or least plausible elements that then receive a certain treatment in the process of belief change. Typically this selection is determined by a preference relation that is independent of the particular revision problem at hand. The second line of defence against my attack on the notion of minimal change in belief change advances the following argument: It is just the upshot of the many AGM-style representation theorems in the literature that rational belief change can be reconstructed as belief change determined by a minimization condition with respect to some underlying doxastic preference relation.<sup>15</sup>

Such preference relations range alternatively over beliefs, sets of beliefs, models or worlds. For the classic accounts of the 1980s, see Gärdenfors (1988, Chapters 3 and 4). And at least in the Grovean possible worlds representation, the minimization procedure has been taken to reflect a form of minimal change: From the models that satisfy the original theory, the agent passes over to the *closest* worlds that satisfy the input sentence.

A first reply to this defence is that contrary to the declarations I cited in Section 2, the AGM postulates for belief revision do not place *any* constraints regarding the preser-

---

worlds in a contraction with respect to  $\phi$ . It says nothing at all about the admission of  $\neg\phi$ -worlds.

<sup>15</sup>In analogy with the tradition of rational choice theory associated with leading economists like Kenneth Arrow and Amartya Sen, one could advance the slogan “rational change is relational change.” In rational choice theory, the corresponding slogan is “rational choice is relational choice.” Belief change thus can be incorporated as a subtheme into the study of the *homo oeconomicus*. As an aside, it may be mentioned that Gärdenfors-Makinson style entrenchment relations can be conceived as revealed preferences in a sense which is common in rational choice theory. See Rott (1996, 1998a).

vation of beliefs in the case where the incoming information is inconsistent with the current theory – the interesting *belief-contravening* case.<sup>16</sup> So if there is no intuition of minimal change or informational economy embodied in the postulates (and this is what I am claiming), it would be strange if we could conjure up such an idea through mathematical representation theorems.

Secondly, I think, it needs to be supported by extra arguments that the canonical preference relations constructed in the proofs of AGM-style representation theorems are more than just technical devices and can indeed be given an interpretation that fits the desired economical meaning. Technically speaking, is of course possible to reverse the underlying preference orderings and consider, for instance, the worlds that are “farthest away” away from the models originally considered possible.

Thirdly, we have seen in Observation 2 that in the case of epistemic entrenchment, the application of the stipulated orderings is not just straightforward minimization.<sup>17</sup> In other constructions like partial meet contraction, safe contraction or possible worlds models, minimization is only one step in a complex procedure of constructing belief changes, and its effects are at least partially neutralized in subsequent steps of this very procedure. There are tacit principles at work here, according to which a believer should respect ties in her underlying preferences and should treat equally objects that she holds in equal regard.<sup>18</sup> And it is again implicit in the construction of AGM-style belief changes that these principles are given priority over principles of minimal change. As a result, these constructions even licence amnesic belief change.

We conclude that the argument to support the idea of minimal change through a rational reconstruction of belief revision in terms of hidden preference relations is not convincing.

### 4.3. Dispositions

The third line of defence of minimal change concedes that the arguments presented in Section 3 are correct if they are viewed as an attack on the particular form of early belief

---

<sup>16</sup>The AGM postulates for belief revision include the above-mentioned postulates of success, logical closure, consistency preservation, as well as a postulate to the effect that it is the content of a new piece of information that matters, not its syntactical formulation. Then AGM have two postulates identifying a revision by a sentence consistent with the prior belief set with the deductive closure of the union of both. (These two postulates are *the only* AGM postulates for revision that lend themselves to a minimal-change interpretation.) Finally, two supplementary postulates of AGM relate potential revisions by varying inputs. For a more extensive interpretation of these postulates, see Rott (1998b).

<sup>17</sup>The Gärdenfors-Makinson recipe for belief contractions based on entrenchments makes use of a disjunction and therefore resists an easy understanding:  $\psi$  is in the contraction of  $K$  with respect to  $\phi$  if and only if  $\psi$  is in  $K$  and  $\phi \vee \psi$  is more entrenched than  $\phi$ .

<sup>18</sup>Compare the discussion of *principles of preference and indifference* in Pagnucco and Rott (1997). It is these principles telling us to go for a compromise if there are various optimal solutions, that make (2) pull in a different direction than the uncompromising minimal change principle (1). See footnote 6.

change models. But this can be blamed on an illegitimate identification of belief states with deductively closed sets of sentences that Alchourrón, Gärdenfors and Makinson seemed to advocate. On a proper understanding, the formal model of a belief state should already encompass the means for changes of beliefs, and should therefore include something like the preference structures discussed in the previous section. Such a move allows us not only to provide a smooth mechanism that can easily cope with iterated changes of belief, but also to get a grip on the elusive notion of minimal change in belief-contravening revisions. If we extend AGM theory and conceive more sophisticated structures (e.g., orderings or selection functions) as representations of doxastic states, we can find natural ways of defining distances between them. The presumption underlying Observation 1, viz., that in order to define minimal changes of belief states one has to compare differences between closed belief sets, is wrong.

Two different but ultimately equivalent ways of implementing this idea are presented by Boutilier (1993, 1996) and Rott (1998c), where the former is based on a possible worlds modelling similar to Grove’s, while the latter is based on epistemic entrenchment relations. Both Boutilier and Rott aim at and furnish an explicit formal understanding of minimal change. The basic idea is that in order to effect a change of the belief state represented by an ordering (of possible worlds or sentences), one should change in the ordering the positions of a uniquely identifiable *minimal*<sup>19</sup> set of pairs, just as it is necessary in order to make the change “successful”. It has turned out, however, that the means to come up with a truly conservative definition of (iterated) belief change has to be dearly purchased. Darwiche and Pearl (1994, 1997) were the first to make this observation by way of intuitive counterexample, and a more general twist is added to it by the following observation made in Rott (1998c).

*Observation 3.* If doxastic states encompass revision-guiding structures (like preference orderings or selection functions), then belief-contravening revisions that obey the maxim of minimum mutilation have unacceptable consequences; they violate a requirement of temporal coherence.

We state this theorem without proof. The gist of the argument (which can be found in the papers mentioned) concerns iterated belief changes and the role of the recency of information. According to the postulate of success, the piece of information that is just coming in (is “most recent”) is maximally appreciated and is therefore invariably accepted. However, when the next revision takes place, the information just taken up (now the “second most recent” belief) turns out to be very weakly entrenched. As Darwiche and Pearl’s (1994, 1997) bird example shows, conservative belief change in this sense can be too dismissive: Revising a *tabula rasa* belief state first by “Fred is a bird”, then by “Fred is red” and finally by “Fred is not a bird” will end up in a

---

<sup>19</sup>Modulo the above-mentioned principle of indifference.

belief state that does not include “Fred is red”, which is certainly a counterintuitive consequence.

In another sense, conservative belief change is too tenacious. Revising the belief state that includes “Barney is either French or Flemish” first by “Barney is not French” and then by “Barney is not Flemish” will preserve the initial disjunction (and even yield that Barney is French), which I think is again contrary to intuition. Incoming information is indeed always accepted, but only at the lowest level possible, so in future cases of conflict it has to take all the blame and is a first candidate for removal.<sup>20</sup> In that respect, therefore, old beliefs are treated with more respect than new beliefs. Conservative belief change as defined by Boutilier and Rott shows no principled stance towards the recency of information and leads to a doxastic behaviour that must be called incoherent with respect to time.

In sum, I do not at all wish to deny that more encompassing representations of belief change are desirable and perhaps even necessary, if we are to deal with the important problem of iterated revisions. However, the most natural suggestions how to reflect the idea of minimal change in such a framework do not lead to a satisfactory solution of the belief revision problem. The argument proposing that the idea of minimal change or informational economy can be supported through an enrichment of the notion of a belief state by structures representing doxastic dispositions is not convincing.

#### 4.4. Truths

Except for the principle of minimum mutilation, little work has been done in the theory of belief revision to account for Quine’s criteria that we mentioned in the introduction. It is particularly irritating that the correspondence of our beliefs with the real world – Quine’s appeal to empirical evidence – is not at all captured by the usual modellings.<sup>21</sup> More than a hundred years ago, William James (1896/1979, p. 24) formulated the main goal of belief fixation and belief change in his famous lecture on “The Will to Believe”: “*We must know the truth; and we must avoid error* – these are our first and great commandments as would-be knowers.” As one of the founding fathers of pragmatism, James is a predecessor to Quine as well as to Isaac Levi, a leading philosopher of belief change who did take over James’s catchword. The charge now against formal accounts

---

<sup>20</sup>The input is accepted, but its negation is considered to be not even moderately surprising.

<sup>21</sup>Symptomatic of a certain lack of interest in *truth* among belief change theorists is the fact that the most important book on *belief* revision has the title “*Knowledge in Flux*” (Gärdenfors 1988), but does not care at all about the truth of our alleged “knowledge.” The same terminological sloppiness is present in the commonly used term ‘*epistemic* entrenchment’ which should really be ‘*doxastic* entrenchment’ – unless one wants to disregard more than two millennia of philosophical tradition (*vide* Plato). That belief revision is finally getting closer to the concept of truth, however, can be seen from the recent work of Kevin Kelly and his colleagues (Kelly, Schulte and Hendricks 1997, Kelly 1998).

of belief change, as well as against my way of capturing minimal change in Principle (1) above, could be that one should worry more about truth. Only true beliefs are valuable, and there is no point in preserving false beliefs. A suitably qualified version of (1) that takes into account this basic concern about truth, so the idea of the fourth defence of minimal change, would not run into the difficulties described above.

The concern with truth is to be welcomed, no doubt, but unfortunately it does not offer an escape from the predicament we have come across. First of all, Observation 1 of course remains valid in a version that replaces belief sets with the sets of *true* beliefs at the respective moments.

*Corollary 4.* Let  $K'$  and  $K''$  be two belief-contravening candidate revisions of a consistent and logically closed belief set  $K$  by a sentence  $\phi$ . If  $K'$  and  $K''$  have different sets of true beliefs, they cannot be set-theoretically compared in terms of the true beliefs on which they differ with  $K$ .

But then, looking at symmetric differences does not seem to be what is called for, since in general we would not mind getting many more true beliefs than we had before the revision. The dyadic notion of minimal change (referring to a relation between old and new belief set) is here affected by the monadic notion of truth (as an aim of the posterior belief set). What we want to make sure is that we minimize the loss of true beliefs. However, this does not lead to a good result either, as is borne out by

*Observation 5.* Let  $K'$  be a belief-contravening candidate revision of a logically closed belief set  $K$  by a sentence  $\phi$ . Then  $K'$  minimizes the set of true beliefs lost from  $K$  (amongst all other candidate revisions of  $K$  by  $\phi$ ) only if  $K'$  is *opinionated* in the sense that it contains either  $\psi$  or  $\neg\psi$ , for every sentence  $\psi$ .

It is clear that the commitment to opinionated theories is an undesirable property, especially since the prior theory may well be undecided about countless matters. Essentially, Observation 5 stands to Observation 3.2 of Alchourrón and Makinson (1982) as Corollary 4 stands to Observation 1. (An independent proof of it is given in the Appendix for the sake of transparency.) It is worth pointing out explicitly that AGM have always been decidedly against the unbridled use of informational economy as it manifests itself in the inflationary behaviour of so-called (maxi-)choice contractions and revisions of logically closed belief sets.<sup>22</sup>

I conclude that the argument to support the idea of minimal change through a relativization to true beliefs does not succeed.

As a little digression, we take down the following, perhaps even more surprising observation.

---

<sup>22</sup>Compare the discussion in Alchourrón and Makinson (1982, pp. 20–21) and in Makinson (1985, pp. 356–359).

*Observation 6.* No belief-contravening candidate revision that does not contain every truth strictly enlarges the set of true beliefs. In particular, even if the prior theory has been false while the new piece of information as well as everything else in the posterior theory is true, the set of truths is not strictly increased.

It turns out that except for God (and perhaps various demons who can jump to the whole truth in a single step) there is no one-way road to the truth.<sup>23</sup> We have to pay a price for every newly acquired truth by trading in other truths. We may hope that these truths are less important or fundamental than the ones we acquire, but this is an issue that cannot be settled in an apriori manner.

## 5. Conclusion

Getting clear about the foundations is not just of theoretical interest, but has a practical effect to it as well: We may expect to learn in which direction to head if we wish to improve our current theories of belief change as well as their application in knowledge-based systems. The main point of this paper is the theories developed by AGM and their followers are *not* oriented toward the principle of informational economy, and we have found no reason why they should. I have called this principle a dogma because many researchers (including the present author) seemed to believe in it and recited it, without actually acting in accordance with it.

We have seen that contrary to many pronouncements of belief change theorists, it is difficult even to spell out what exactly is meant by the idea of informational economy or minimal change. Theories of belief change in the tradition of Alchourrón, Gärdenfors and Makinson do not align easily with the idea of minimal change, and they are certainly not at all centered on this idea. I have tried to show that two of the most natural ways of fleshing out the idea do not achieve what they are supposed to achieve, and that four attempts to defend minimal change against the attack fail.

I should like to stress, however, that I have only been concerned with principles that are supposed to give a *description* of what motivates an important class of existing normative theories about belief change. I have refrained from putting forward any *normative* thesis about belief change myself, neither against nor in favour of conservatism. This is an entirely different undertaking.<sup>24</sup>

---

<sup>23</sup>James (1896/1979, p. 24) seems to have felt that such problems may arise: "... it may indeed happen that when we believe the truth *A*, we escape as an incidental consequence from believing the falsehood *B* .... We may in escaping *B* fall into believing other falsehoods, *C* or *D*, just as bad as *B* ..."

<sup>24</sup>See especially Harman (1986). A good critical discussion of normative conservatism in contemporary epistemology is given in Christensen (1994).

At the end of the paper, another point has come up. Little work – perhaps no work at all – has been done that reflects Quine’s criteria (a) and (b) in the theory of belief revision. In his joint book with J.S. Ullian, *The Web of Belief* (1978), Quine has added more virtues that good theories should have: modesty, generality, refutability, and precision. Again, belief revision as studied so far has little to offer to reflect the quest for these intuitive desiderata. Except for the issue of conservatism, Quine’s list is one of theory *choice* rather than theory *change* in that it lists properties that a good posterior theory should have, independently of the properties of the prior theory. It is a strange coincidence that the philosophy of science has focussed on the monadic (non-relational) features of theory choice, while philosophical logic has emphasized the dyadic (relational) features of theory change. I believe that it is time for researchers in both fields to overcome this separation and work together on a more comprehensive picture.

With the present paper, however, I first of all hope to draw attention to the fact that it is not appropriate to exclusively focus on informational economy even when talking on nothing but the restricted, AGM-style modellings of belief change. There are various coherence criteria that find expression in formal “rationality postulates”: inferential coherence (consistency and closure), dispositional coherence (a kind of semantic representability syntactically encoded in so-called “supplementary” rationality postulates – which are what I would call the hallmark of AGM theory), as well as temporal coherence (a principled appreciation of recency of information in iterated belief change). These criteria substantially restrict the dominion of informational economy in formal belief change theories. Having understood this, we may expect that the picture will change again enormously when we shall actually be going to live up to the Quinean epistemological virtues. It is time to liberate ourselves from the constrictions of informational economy and try out other norms for our ethics of belief.

## Acknowledgments

I warmly thank Peter Gärdenfors for all his generous and undogmatic help over many years. Happy birthday to you, Peter!

An early version of this paper was presented at the Third Conference on Logic and the Foundations of the Theory of Games and Decisions (LOFT3, Torino, December 1998). I am grateful to Isaac Levi, David Makinson, Oliver Schulte, and the participants of LOFT3 for a number of very valuable comments on various versions of this paper.

## Appendix: Proofs.

*Observation 1.* No two distinct belief-contravening candidate revisions of a consistent

and logically closed belief set by a sentence  $\phi$  can be set-theoretically compared in terms of the sets of beliefs on which they differ with the original belief set.

*Proof of Observation 1.* Let  $\neg\phi$  be in a consistent belief set  $K$ , and let  $K'$  and  $K''$  be two distinct candidate revisions of  $K$  with respect to  $\phi$ . By the postulates of Closure and Success, we know that both  $K'$  and  $K''$  are logically closed and contain  $\phi$ .

We want to show that there is a sentence which is in  $K\Delta K'$  but not in  $K\Delta K''$  and there is also a sentence that is in  $K\Delta K''$  but not in  $K\Delta K'$ .

Since  $K \neq K'$  there is either a sentence in  $K' - K''$  or there is a sentence in  $K'' - K'$ . Without loss of generality, assume the latter, and take some sentence  $\psi$  that is contained in  $K''$  but not in  $K'$ . Then, by Closure,  $\neg\phi \vee \psi$  is in  $K$  and in  $K''$  but not in  $K'$ . Hence  $\neg\phi \vee \psi$  is in  $K\Delta K'$  but not in  $K\Delta K''$ . On the other hand, again by Closure,  $\phi \wedge \psi$  is in  $K''$  but neither in  $K$  nor in  $K'$  (here we also use the consistency of  $K$ ). Hence  $\phi \wedge \psi$  is in  $K\Delta K''$  but not in  $K\Delta K'$ . We conclude that  $K\Delta K'$  and  $K\Delta K''$  are not related by subset inclusion.  $\square$

*Observation 2.* Suppose we want to revise a belief set by a sentence  $\phi$  and find two elements of the belief set that non-redundantly entail the negation of  $\phi$ . Then it may well be rational, according to the standard belief revision constructions, to restore consistency by removing the more entrenched and retain the less entrenched belief. In fact, such a situation can always be identified in an anamnestic revision by a consistent and moderately surprising sentence.

*Proof of Observation 2.* First we give a simple example of a situation with the characteristics depicted above. Consider  $K = Cn(\neg p, q)$  and the following sequence of sentences of decreasing logical strength:  $\perp$ ,  $\neg p \wedge q$ ,  $\neg p$ ,  $p \supset q$  and  $\top$ . Define the ordering  $\leq$  by putting  $\phi \leq \psi$  iff every sentence in the sequence which implies  $\phi$  also implies  $\psi$ .

An illustration of this situation in terms of Grovean systems of spheres (cf. Grove 1988) is given in Figure 1 in the main text above.

It is easy to check that this defines a standard entrenchment relation in the sense of Gärdenfors and Makinson (1988). By the construction recipe for entrenchment-based revisions, we get  $K * p = Cn(\{\phi \in K : \neg p < \phi\} \cup \{p\})$  which equals  $Cn(p, q)$ . Evidently  $q$  and  $q \supset \neg p$  form a non-redundant set that entails  $\neg p$  (“a minimal culprit set for  $\neg p$ ”), and  $q$  is strictly less entrenched than  $q \supset \neg p$ . And yet  $q$  is, but  $q \supset \neg p$  is not, contained in  $K * p$ .

For the general part of Observation 2, let  $\phi$  be a consistent and moderately surprising sentence, and let  $K * \phi$  be unequal to  $Cn(\phi)$ . The former means that there is an  $\alpha$  in  $K$  such that  $\alpha < \neg\phi$ , where  $<$  is the entrenchment relation associated with  $K$ . Take such



an  $\alpha$ , and let  $\beta$  be an element of  $K * \phi$  which is not implied by  $\phi$ . We need to find two beliefs  $\psi$  and  $\chi$  which jointly, but not individually imply  $\neg\phi$ , and are such that  $\psi < \chi$  but  $\psi$  is maintained while  $\chi$  is abandoned in  $K * \phi$ . Consider

$$\begin{aligned}\psi &:= (\neg\phi \wedge \alpha) \vee (\phi \wedge \beta) \\ \chi &:= \neg\phi \vee \neg\beta\end{aligned}$$

We check that these sentences  $\psi$  and  $\chi$  indeed have the desired properties (for an illustration in terms of Grove spheres, compare Figure 2).

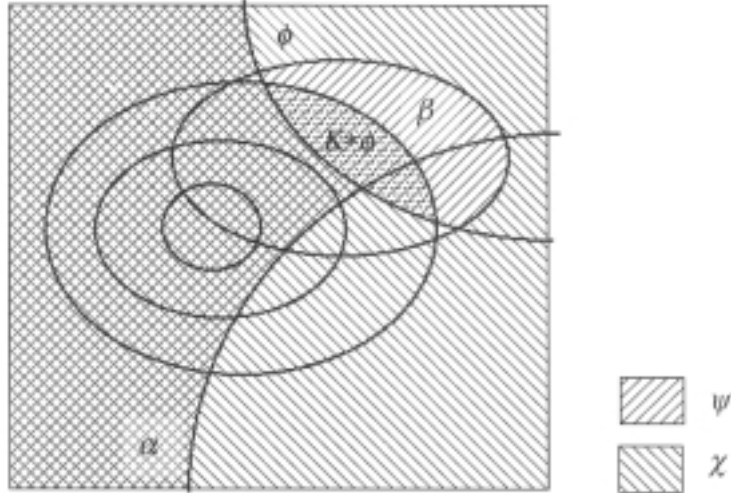


Figure 2: Illustration of the proof of the general part of Obs. 2

First,  $\psi \wedge \chi$  clearly implies  $\neg\phi$ . Second,  $\psi$  alone does not imply  $\neg\phi$ , since  $\phi$  and  $\beta$  are both elements of the consistent set  $K * \phi$  (applying AGM's postulate that consistency can never be occasioned by a consistent input  $\phi$ ), and thus their conjunction is also consistent. Third,  $\chi$  alone does not imply  $\neg\phi$ , since  $\phi$  does not imply  $\beta$  and therefore  $\neg\beta$  does not imply  $\neg\phi$ . Fourth,  $\psi$  is in  $K * \phi$  since both  $\phi$  and  $\beta$  are in  $K * \phi$ , and  $K * \phi$  is closed under logical consequence (using the AGM postulates of Success and Closure). Fifth,  $\chi$  is not in  $K * \phi$ , since  $\phi \wedge \beta$  is in  $K * \phi$  and  $K * \phi$  is consistent (again assuming Success and Consistency Preservation). Finally, it remains to show that  $\psi < \chi$ . (In the following we refer to the convenient axiomatization of entrenchment relations given in Rott 1992b.) From the irreflexivity of  $<$ , we get that  $\alpha \not< \alpha$ . Since  $\neg\phi \wedge ((\neg\phi \wedge \alpha) \vee (\phi \wedge \beta))$

implies  $\alpha$ , we get by Continuing Up that  $\alpha \not\prec \neg\phi \wedge ((\neg\phi \wedge \alpha) \vee (\phi \wedge \beta))$ . From this and  $\alpha < \neg\phi$ , we get by Conjunction Up that  $\alpha \not\prec (\neg\phi \wedge \alpha) \vee (\phi \wedge \beta)$ . Making use of the fact that  $\alpha < \neg\phi$  once more, we deduce from this with the help of Virtual Connectivity (a property characteristic of Gärdenfors-Makinson entrenchments) that  $(\neg\phi \wedge \alpha) \vee (\phi \wedge \beta) < \neg\phi$ . By Continuing Up, we finally get  $(\neg\phi \wedge \alpha) \vee (\phi \wedge \beta) < \neg\phi \vee \neg\beta$ , that is,  $\psi < \chi$ , as desired. (Notice that this proof does not depend on a particular construction recipe for entrenchment-based belief changes.)  $\square$

*Observation 5.* Let  $K'$  be a belief-contravening candidate revision of a logically closed belief set  $K$  by a sentence  $\phi$ . Then  $K'$  minimizes the set of true beliefs lost from  $K$  (amongst all other candidate revisions of  $K$  by  $\phi$ ) only if  $K'$  is opinionated in the sense that it contains either  $\psi$  or  $\neg\psi$ , for every sentence  $\psi$ .

*Proof of Observation 5.* Let the belief set  $K$  be consistent and logically closed, let  $\neg\phi$  in  $K$ , and let  $K'$  be a candidate revision containing  $\phi$  which is not opinionated. We show that  $K'$  does not minimize the set of lost true beliefs from  $K$ .

Assume first that  $\phi$  is true. Then the set  $T$  of all true sentences is a candidate revision containing  $\phi$ . Take a true sentence  $\psi$  that is not contained in  $K'$ ; there is such a  $\psi$  since  $K'$  is not opinionated. Then, by logical closure,  $\neg\phi \vee \psi$  is a true sentence that was contained in the prior belief set  $K$  and is lost in the candidate belief set  $K'$ . But  $\neg\phi \vee \psi$  is in  $T$ , and  $T$  loses no true belief from  $K$ , so  $T$  actually loses less true beliefs from  $K$  than  $K'$ . Thus  $K'$  does not minimize the set of lost true beliefs from  $K$ .

Assume secondly that  $\phi$  is false. Then  $\neg\phi$  is true. Extend  $K'$  to an opinionated set  $K''$ , and take some  $\psi$  from  $K'' - K'$ . Then, by logical closure,  $\neg\phi \vee \psi$  is a true sentence from  $K$  which is lost in  $K'$  but not in  $K''$ . Since  $K'$  is a subset of  $K''$ ,  $K''$  loses no true sentence from  $K$  that  $K'$  doesn't lose. So  $K''$  actually loses less true beliefs from  $K$  than  $K'$ . Thus  $K'$  does not minimize the set of lost true beliefs from  $K$ .  $\square$

*Observation 6.* No belief-contravening candidate revision that does not contain every truth strictly enlarges the set of true beliefs. In particular, even if the prior theory has been false while the new piece of information as well as everything else in the posterior theory is true, the set of truths is not strictly increased.

*Proof of Observation 6.* Let  $\neg\phi$  be in the belief set  $K$ , and let  $K'$  be a potential revision of  $K$  with respect to  $\phi$ . By the postulates of Closure and Success, we know that  $K'$  is logically closed and contains  $\phi$ .

Again we let  $K_t$  and  $K'_t$  be the set of true beliefs in  $K$  and  $K'$  respectively. What we want to show that  $K'_t$  is not a strict superset of  $K_t$ , i.e., that there is a true sentence which is in  $K$  but is lost in  $K'$ .

Take some true sentence  $\psi$  which is not in  $K'$ . Such a sentence exists since  $K'$  was assumed not to be omniscient. Now consider  $\neg\phi \vee \psi$ . This sentence is true, since  $\psi$  is true, and it is in  $K$ , since  $\neg\phi$  is in  $K$  and  $K$  is logically closed. However,  $\neg\phi \vee \psi$  is not in  $K'$  since  $\phi$  is in  $K'$  and  $\psi$  is not in  $K'$ . Thus we have found a truth that has been lost.

The second part of Observation 6 follows from the first.  $\square$

## References

- Alchourrón, Carlos, Peter Gärdenfors and David Makinson: 1985, 'On the Logic of Theory Change: Partial Meet Contraction Functions and Their Associated Revision Functions', *Journal of Symbolic Logic* **50**, 510–530.
- Alchourrón, Carlos, and David Makinson: 1982, 'On the Logic of Theory Change: Contraction Functions and Their Associated Revision Functions', *Theoria* **48**, 14–37.
- Alchourrón, Carlos, and David Makinson: 1985, 'On the Logic of Theory Change: Safe Contraction', *Studia Logica* **44**, 405–422.
- Boutilier, Craig: 1993, 'Revision Sequences and Nested Conditionals', in R. Bajcsy (ed.), *IJCAI-93 – Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 519–525.
- Boutilier, Craig: 1996, 'Iterated Revision and Minimal Change of Conditional Beliefs', *Journal of Philosophical Logic* **25**, 263–305.
- Christensen, David: 1994, 'Conservatism in Epistemology', *Noûs* **28**, 69–89.
- Darwiche, Adnan, and Judea Pearl: 1994, 'On the Logic of Iterated Belief Revision', in Ronald Fagin, ed., *TARK'94 – Proceedings of the Fifth Conference on Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, Pacific Grove, Cal., pp. 5–23.
- Darwiche, Adnan, and Judea Pearl: 1997, 'On the Logic of Iterated Belief Revision', *Artificial Intelligence* **89**, 1–29.
- Fuhrmann, André: 1997, *An Essay on Contraction*, CSLI Publications, Stanford.
- Gärdenfors, Peter: 1988, *Knowledge in Flux. Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, Cambridge, Mass.
- Gärdenfors, Peter: 1992, 'Belief Revision: An Introduction', in P.G. (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 1–28.
- Gärdenfors, Peter, and David Makinson: 1988, 'Revisions of Knowledge Systems Using Epistemic Entrenchment', in Moshe Vardi (ed.), *TARK'88 – Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Morgan Kaufmann, Los Altos, pp. 83–95.
- Gärdenfors, Peter, and Hans Rott: 1995, 'Belief revision', in D. M. Gabbay, C. J. Hogger, and J. A. Robinson (eds.), *Handbook of Logic in Artificial Intelligence and Logic Programming Volume IV: Epistemic and Temporal Reasoning*, Oxford University Press, pp. 35–132.
- Grove, Adam: 1988, 'Two Modellings for Theory Change', *Journal of Philosophical Logic* **17**, 157–170.

- Hansson, Sven O.: 1999a, *A Textbook of Belief Dynamics: Theory Change and Database Updating*, Dordrecht: Kluwer Academic Publishers.
- Hansson, Sven O. (ed.): 1999b, Special Issue on Non-Prioritized Belief Revision, *Theoria* **63**, 1–134.
- Harman, Gilbert: 1986, *Change in View*, Bradford Books, MIT Press, Cambridge, Mass.
- James, William: 1979, ‘The Will to Believe’, in W.J., *The Will to Believe and Other Essays in Popular Philosophy*, The Works of William James, Vol. 6, edited by F.H. Burkhardt, F. Bowers and I.K. Skrupskelis, Cambridge, Mass., and London, Harvard University Press, pp. 13–33. (Orig. 1896)
- Kelly, Kevin: 1998, ‘The Learning Power of Belief Revision’, in Itzhak Gilboa (ed.), *Theoretical Aspects of Rationality and Knowledge. Proceedings of the Seventh Conference (TARK 1998, Evanston/Illinois)*, San Francisco, Morgan Kaufmann, pp. 235–248.
- Kelly, Kevin, Oliver Schulte, and Vincent Hendricks: 1997, ‘Reliable Belief Revision’, in Maria Luisa Dalla Chiara et al. (eds.), *Logic and Scientific Methods – Proceedings of the 10th International Congress of Logic, Methodology and Philosophy of Science*, Dordrecht, Kluwer.
- Makinson, David: 1985, ‘How to Give it Up: A Survey of Some Formal Aspects of the Logic of Theory Change’, *Synthese* **62**, 347–363.
- Miller, David: 1974, ‘Popper’s Qualitative Theory of Verisimilitude’, *British Journal for the Philosophy of Science* **25**, 166–177.
- Pagnucco, Maurice, and Hans Rott: 1997, ‘Severe Withdrawal (and Recovery)’, to appear in the *Journal of Philosophical Logic*.
- Quine, Willard V.O.: 1951, ‘Two dogmas of empiricism’, *Philosophical Review* **60**, 20–43. Reprinted in W.V.O. Quine, *From a Logical Point of View*, Harvard University Press, Cambridge, Mass, 1953, pp. 20–46.
- Quine, Willard V.O., and Joseph S. Ullian: 1978, *The Web of Belief*, Random House, New York, second edition.
- Rott, Hans: 1991, ‘Two Methods of Constructing Contractions and Revisions of Knowledge Systems’, *Journal of Philosophical Logic* **20**, 149–173.
- Rott, Hans: 1992, ‘On the Logic of Theory Change: More Maps Between Different Kinds of Contraction Function’, in Peter Gärdenfors (ed.), *Belief Revision*, Cambridge University Press, Cambridge, pp. 122–141.
- Rott, Hans: 1992, ‘Preferential Belief Change Using Generalized Epistemic Entrenchment’, *Journal of Logic, Language and Information* **1**, 45–78.
- Rott, Hans: 1996, *Making Up One’s Mind: Foundations, Coherence, Nonmonotonicity*, Habilitationsschrift, Department of Philosophy, University of Konstanz, October 1996. To appear under the title ‘Change, Choice and Inference’ with Oxford University Press.
- Rott, Hans: 1998a, ‘Logic and Choice’, in Itzhak Gilboa (ed.), *Theoretical Aspects of Rationality and Knowledge. Proceedings of the Seventh Conference (TARK 1998, Evanston/Illinois)*, San Francisco, Morgan Kaufmann, pp. 235–248.
- Rott, Hans: 1998b, ‘Coherence and Conservatism in the Dynamics of Belief. Part I: Finding the right framework’, to appear in *Erkenntnis*, special issue on Coherence and Belief Revision, eds. V. Halbach and E. Olsson.

- Rott, Hans: 1998c, 'Coherence and Conservatism in the Dynamics of Belief. Part II: Iterated belief change without dispositional coherence', manuscript, ILLC, University of Amsterdam, May 1998.
- Tichý, Pavel: 1974, 'On Popper's Definition of Verisimilitude', *British Journal of the Philosophy of Science* **25**, 155–160.