THE JOURNAL OF SYMBOLIC LOGIC
Volume 71. Number 2. June 2006

ON THE DEGENERACY OF THE FULL *AGM*-THEORY OF THEORY-REVISION

NEIL TENNANT

Abstract. A general method is provided whereby bizarre revisions of consistent theories with respect to contingent sentences that they refute can be delivered by revision-functions satisfying both the basic and the supplementary postulates of the AGM-theory of theory-revision.

§1. Introduction.

1.1. Preliminaries. Fix some language L, and a deducibility relation $\vdash$ on L. All that matters is that $\vdash$ should obey the usual structural rules of reflexivity, dilution and cut, and should afford the primitive inferential moves corresponding to introducing dominant occurrences of $\neg$ (negation), $\wedge$ (conjunction) and $\rightarrow$ (implication) on the right and on the left of sequents. Note that intuitionistic logic affords all these inferential moves, as does classical logic. We use the symbol $\bot$ for absurdity. (So a set K of sentences is inconsistent just in case $K \vdash \bot$.) We shall use A, B, C and D for sentences, and use K, J and H for sets of sentences. Occasionally we shall make use also of φ as a sentence-parameter and Θ as a parameter for sets of sentences.

DEFINITION 1. *The* logical closure $[K]$ *of a set* K *of sentences* (*in* L) *is* $\{\varphi \in L \mid K \vdash \varphi\}$. *We write* $[K, J]$ *for* $[K \cup J]$ *and write* K, A *for* $K \cup \{A\}$.

DEFINITION 2. $K \vdash J$ *will mean that for every sentence* φ, *if* $J \vdash \varphi$ *then* $K \vdash \varphi$. *If* J *is a proper subtheory of* K *and* $K \vdash A$, *then we say that* J permits recovery of K via A *just in case* $J, A \vdash K$.

DEFINITION 3. A theory *in L is a logically closed set of sentences in L, i.e., a set K of sentences in L such that* $K = [K]$; *equivalently, such that for every sentence* $\varphi \in L$, *if* $K \vdash \varphi$ *then* $\varphi \in K$.

Any set of beliefs (presumed to be held by a rational agent) will here be modelled, as usual, as a consistent theory (in some language L). This modelling assumption is prevalent among belief-revision theorists. Henceforth reference to the language L will be suppressed

Received June 28, 2005.

Thanks are owed to Sven Ove Hansson, David Makinson, Daniel Osherson, Mark Ryan, Timothy Smiley, Wolfgang Spohn and Gabriel Uzquiano-Cruz for their comments on earlier drafts. Suggestions and criticisms from anonymous referees have also led to improvements in the exposition. Special thanks are owed to Harvey Friedman, who in response to Theorem 4 raised the question whether the AGM postulates would allow theories to be bizarrely revisable in infinitely various ways. That put the author on a path of investigation that arrived at Theorem 5, which provides an affirmative answer to the question. The author is solely responsible for any defects that remain.

0022-4812/06/7102-0017/$2.60

Belief revision is understood here as the process of revising a consistent theory K with respect to a contingent sentence A that is inconsistent with K. It is assumed that a rational agent would have a method for obtaining a unique result. That method would be encoded by a revision-function $*$, which would depend on the agent in question. The result of such revision is denoted as $K * A$.

There are two other theory-change operations to note besides that of revision: these are the operations of *expansion* and of *contraction*. Intuitively, one *expands* one's belief-set K with respect to A by adding A to K (and then taking the logical closure of the result). So the *expansion* $K + A$ (of K with respect to A) is $[K, A]$. Here, moreover, we assume that A is consistent with K. In order to *contract* K with respect to A—that is, to *get rid of* A—it can be assumed that A is already in K, and is not itself a logical truth (for we do not wish ever to get rid of any of those). The *contraction* (of K with respect to A) is denoted as $K - A$.

In order to *revise* K with respect A—that is, to adopt the new belief A even though it *conflicts* with one's current beliefs K—it may be assumed that A is not itself a logical falsehood, but is inconsistent with K. If, however, A is *consistent* with one's present belief-set K, then the operation of adopting A as a new belief is simply that of *expanding* K with respect to A: 'revision' would, in such circumstances, be a misnomer. For more on the how the domains of definition of expansion, contraction and revision should be construed as trichotomizing the space of logical possibilities, see Tennant [to appear].

The three operations of expansion, contraction and revision all feature in the well-known 'Levi identity':

$$K * A = (K - \neg A) + A.$$

Definition 4. *We say 'K is ripe for revision with respect to A' if and only if the following three conditions hold*: $K \nvdash \bot$; $A \nvdash \bot$; $K, A \vdash \bot$.

Definition 5. *Let the theory K imply the contingent sentence A. A* maxichoice contraction *of K with respect to A is a maximal non-(A-implying) subset of K: and $K \perp A$ is the set of all such subsets of K. The* full meet contraction *of K with respect to A is* $\bigcap(K \perp A)$.

Definition 6. *Let K be ripe for revision with respect to A. Any maxichoicecontraction J of K with respect to $\neg A$ induces the corresponding* maxichoice revision $[J, A]$ *of K with respect to A. The full meet contraction* $\bigcap(K \perp \neg A)$ *of K with respect to $\neg A$ induces the corresponding* full meet revision $[\bigcap(K \perp \neg A), A]$ *of K with respect to A.*

Suppose K is ripe for revision with respect to A. Suppose our rational agent is using the revision function $*$. Then the revised theory $K * A$ should contain A. This is the requirement of SUCCESS. $K * A$ should also contain 'as much of K as is possible'. This is the requirement of 'informational economy'. It is also known as the *maxim of minimal mutilation*. A seldom-remarked further requirement is that $K * A$ should contain 'as little else as is possible'. One might call this the *maxim of minimal bloating*.[1] It says that the process of revision should not, by itself, lead one to adopt any new beliefs that have no business being in the revised theory.

[1]The author uses the term 'bloating', rather than the more neutral term 'inflation', in order to connote that it is something very much to be avoided. After all, belief-revision theorists already use the term 'mutilation', rather than the more neutral-sounding 'deflation', in order to convey a similar connotation at the other extreme.

It is important to note, for the purposes of assessing the significance of the results in this paper, that the pre-formal notions of minimal mutilation and minimal bloating can remain *unanalyzed.* No formal explication of these pre-formal notions is called for. We do not need to take any stand on the question whether these pre-formal notions admit of any satisfactory formal explication at all. Nor need we take any view as to which of any two competing explications might be the better one. The results below enjoy their significance independently of any answer to such questions. This is because the results can be framed in terms of the *pre-formal* notions.

At the appropriate juncture below, when the dialectical setting is clearer, we shall re-visit the foregoing point in order to emphasize it. To foreshadow: what this paper shows is how to accommodate, in a revision-function satisfying all the so-called *AGM* postulates for revision (for which, see §2 below), both bloatings and mutilations of any kind that the reader might have in mind—not just those mutilations produced by full-meet revisions or those bloatings produced by maxichoice revisions (for which, see §1.2 below). To make matters worse: the bizarre revisions that the *AGM* postulates admit can combine both extreme mutilation and extreme bloating (however the reader wishes to construe these). And to make matters even worse: such bizarre revisions can take place, as it were, 'infinitely-where'.

1.2. Historical background. The *AGM* literature to date reveals the possibility of only two special cases of degeneracy. Moreover, these possibilities have so far been revealed to exist only in isolation, rather than in combination. Let us assume, in order to explain these two cases, that K is ripe for revision with respect to A.

Alchourrón and Makinson [1982] made the following two observations. First, every maxichoice revision of K with respect to A is a complete theory.[2] Hence, for incomplete theories K and contingent sentences A of minor consequence, maxichoice revisions of K with respect to A will be (maximally) bloating. Secondly, the full meet revision of any theory K with respect to A is just $[A]$ itself.[3] Hence, for non-trivial K and incomplete $[A]$, the full meet revision of K with respect to A will be (maximally) mutilating.

The first of these points was reiterated by Alchourrón, Gärdenfors and Makinson [1985], at p. 511, and was used to motivate the introduction of the notion of *partial meet* contractions and their associated revisions. Immediately after raising the problem that maxichoice revisions are complete, they wrote as follows:[4]

> The "inflation properties" that ensue from applying the maxichoice operations bring out the interest of looking at other formal operations that yield smaller sets as values. In this paper, we will start out from the assumption that there is a selection function γ that picks out a class of the "most important" maximal subsets of K that fail to imply A. The contraction $K \dot{-} A$ is then defined as the intersection of all the maximal subsets selected by γ. Functions defined in this way will be called *partial meet contraction* functions, and their corresponding revision functions will be called *partial meet revision* functions.

[2]Observation 3.2, p. 21.

[3]Observation 2.2, p. 19.

[4]*Loc. cit.*, p. 511. Notation for theories and sentences has been changed to that of this paper.

So: in the case of the contraction of K with respect to A. meets would be taken of *some several.* but *not necessarily all.* maximally non-(A-implying) subsets of K. This was clearly intended to avoid both maxichoice contraction (where a single such subset of K would be taken. so the meet would be degenerate) and full meet contraction (where all such subsets of K would be taken. and the meet would accordingly be inclusion-minimal). Although in the above quote we see motivating mention only of the problem of bloating ('inflation') posed by maxichoice revisions. it is clear that the authors desired also to avoid the problem of mutilation posed by full meet revision.

It is also evident that the move to partial meet revisions was being thought of as a way of avoiding the respective kinds of degeneracy at the two extremes. Recall that single members of $(K \perp \neg A)$ give rise to maxichoice revisions of K with respect to A. with the concomitant possibility of bloating (but not necessarily of mutilation): while the full set $(K \perp \neg A)$ gives rise to the full meet revision of K with respect to A. with the concomitant possibility of (maximal) mutilation but not of bloating. Alchourrón. Gärdenfors and Makinson moved. in their 1985 paper. from the main motivating case of bloating by maxichoice revisions (which they called 'inflation') directly to their new proposal to consider partial meet revisions.

Given their professed motivation. it is significant that they did not raise or address the worry that either kind of degeneracy (the bloating inflicted by maxichoice revisions. or the mutilation inflicted by the full meet revision) might persist with infinitely many choices of non-empty subsets of $(K \perp \neg A)$ that could be used to construct partial meet revisions of K with respect to A. They did not anticipate that the two kinds of degeneracy might arise for partial meet revisions. either in isolation or—which would be even worse—in combination. The operational assumption that the reader would be justified in imputing to them was that the maxichoice and full-meet cases could be viewed as lying at opposite ends of what could now be presumed to be a spectrum with a reasonable 'middle range' throughout. so to speak. That was the spectrum of (transitively relational) partial-meet contraction-functions. The corresponding revision-functions would then be obtained via the Levi identity.

As evidence for the quietist view that this reading attributes to AGM-theory. note that Gärdenfors [1988] claimed (at p. 82) that because the basic and supplementary AGM-postulates for contraction had been 'independently motivated'. the partial meet representation theorem for the contraction postulates

> gives us a strong reason to focus on transitively relational partial meet contraction functions as an *ideal* representation of the intuitive process of contraction. [Emphasis added.]

Likewise. then. the corresponding partial meet representation theorem for the revision postulates should (by straightforward adaptation of the foregoing quote) '[give] us a strong reason to focus on transitively relational partial meet [revision] functions as an *ideal* representation of the intuitive process of [revision].'

The thought seemed to be. therefore. that by avoiding the two extremes and sticking to the middle range. the postulates would now permit only rationally admissible revisions to take place. and the choice among the various competing belief-revision functions in this middle range would be a largely pragmatic matter. depending on varying preferences among rational agents. If this were *not* the thought. then it is inexplicable how potential worries about both bloating and mutilation, as legitimately provoked by maxichoice revision and full-meet revision respectively, did not continue to be explicitly and actively entertained by *AGM*-theorists with regard to partial-meet revisions in the newly provided middle range. Instead, it appeared to be an article of faith that partial meet revisions of K with respect to A that were constructed in this middle range by means of non-empty, proper, subsets of $(K \perp \neg A)$ would avoid the two extremes and their respective kinds of degeneracy. They would involve neither mutilation nor bloating. Such was the reassurance that the method of partial meets seemed to provide.

The results of this paper show that even the middle range is fraught, in infinitely many places, with degeneracy. Completely bizarre revisions are permitted throughout this middle range. We prove a result of the form '*any* kind of bizarre-looking revision is sanctioned by the full set of *AGM* postulates'. One can say to the reader, in effect: 'Choose any bizarre revisions you like, by whatever criteria of bizarreness you care to employ; they can be yielded by a revision function that satisfies all the *AGM* postulates, both basic and supplementary.' Note that maxichoice revisions can be bloating (but no one claimed they were mutilating); and full meet revisions can be mutilating (and certainly they are not bloating). So it is worth emphasizing here that when the reader is invited, as just described, to choose any bizarre revisions she likes, the present author means that she may choose (individual) revisions that are *both bloating and mutilating*.

Nothing like this has hitherto been anticipated by *AGM* theorists. If it *had* been anticipated—let alone proved—then there should have been considerable consternation about the adequacy of the *AGM* postulates as an account of rational belief-revision. For the problem is not just that one can find some revision functions (satisfying the postulates) that are bloating and that one can find yet other revision functions (satisfying the postulates) that are mutilating. The problem is much worse than this!—one can find revision functions (satisfying the postulates) that are both bloating and mutilating. The text quoted above from Alchourrón, Gärdenfors, and Makinson [1985] continues as follows:

> It will be shown that they [partial meet contraction and revision functions —NT] satisfy Gärdenfors' postulates, and indeed provide a representation theorem for those postulates. When constrained in suitable ways, by relations or, more restrictedly, by transitive relations, they also satisfy his "supplementary postulates", and provide another representation theorem for the entire collection of "basic" plus "supplementary" postulates.

The results to be presented below reveal the 'entire collection of "basic" plus "supplementary" postulates' for belief-revision functions to be surprisingly lax. Whether this degree of laxity is acceptable (on the part of a theory aspiring to characterize rational belief-revision) is a more philosophical and methodological question into which we do not, in this paper, inquire. Our concerns here are purely logical.

The upshot of our investigations might be put this way: what the celebrated representation theorem was supposed to establish was a conceptual bridge between a postulationally constrained notion and a mathematically constructed one. But all that this representation theorem shows is that a totally inadequately constrained notion of revision implicitly defined by the chosen postulates is matched by a totally inadequately constrained notion of revision constructed by means of partial meets. So much the worse. then. for the representation theorem. Bizarre and undesirable revisions are still allowed to reside at both ends of the bridge.

1.3. Comments on our methods of proof. A note on the style of argumentation employed here is in order. The proofs below focus on the deducibility relation. But they exploit only rather weak properties of that relation. In particular. our results would apply to an object-language whose logic was intuitionistic. All the results could be given algebraic proofs instead. by means of 'semantic diagrams'. interpreting theories as sets of possible worlds. The latter method. however. can easily commit one to a classical logic for the object-language. unless one is extremely circumspect. (Moreover. the author has written out the algebraic proofs for all the results of this paper. only to find that there is no significant reduction to be had in the overall length of proofs.)

It is desirable to have an account of theory-revision that would deal with intuitionistic theories as well as classical ones.[5] To be sure. there are occasional points at which we apply constructive dilemma or classical *reductio* in our reasoning. But that is a use of classical logic in the metalanguage. not the object-language. Moreover. such uses of classical reasoning in the metalanguage are applied to statements of deducibility in the object-language. So these uses would actually be acceptable to the intuitionist provided that the deducibility relation in the object-language was decidable—which of course is the case for both intuitionistic and classical propositional logic. The meta-reasoning is therefore *strictly* classical only in the case where the logic of the object-language is undecidable (such as would be the case with intuitionistic first-order logic with at least two monadic predicates. and with classical first-order logic with at least one dyadic predicate). By focusing on deducibility. however. rather than using semantic diagrams. we obtain results with the ready assurance that they are applicable not only to classically closed theories. but also to theories closed under intuitionistic. even if not classical. logic.

§2. The *AGM*-theory of revision.

2.1. The basic postulates for revision. We focus in this paper on the operation of revision. bearing in mind that revision is intrinsically related to the other two operations of expansion and contraction.

The *AGM*-postulates governing belief-revision functions can be found in Alchourrón. Gärdenfors and Makinson [1985] at pp. 513 and 515. and (in slightly altered form) in Gärdenfors [1988]. pp. 54–56. For definiteness. we focus on the earlier paper. It lays down six 'basic' postulates (K*1)–(K*6). For present purposes. we may assume without loss of generality that K is ripe for revision with respect to A (this being the principal case of theoretical interest).

(K*1) $K * A$ is always a theory.

(K*2) A is in $K * A$.

(K*3) If A is consistent with K. then $K * A$ is $[K. A]$. (This postulate becomes unnecessary if. as here. $K * A$ is defined only when K refutes A. which is required in order for K to be ripe for revision with respect to A.)

[5]See Tennant [2005] for a discussion of the problems involved when generalizing the *AGM* account of contraction in order to deal with intuitionistic theories.

(K*4) If A is consistent. then so is $K * A$.
(K*5) If A and B are logically equivalent. then $K * A = K * B$.
(K*6) $(K * A) \cap K$ is $K - \neg A$. the *contraction* of K with respect to $\neg A$.

Postulates (K*1) and (K*4) ensure that revisions are consistent theories. Postulate (K*5) ensures that revisions are sensitive only to the logical content of A. Postulate (K*2) ensures that revisions are successful.

Postulate (K*6) addresses the problem of mutilation only indirectly. via contraction. One of the postulates governing contraction-functions is that of RECOVERY:

(K−6) Every theorem of K is a theorem of $[(K - A). A]$.

(K−6) may also be expressed by saying that the contraction $K - A$ permits recovery of K via A (recall Definition 2). (K−6) is *AGM*'s formal attempt to capture the requirement that contractions should be minimally mutilating. So we see that *AGM* seeks to impose the corresponding requirement on revisions via (K*6). which identifies the common part of K and $(K * A)$ as the contraction of K with respect to $\neg A$. It is an easy logical exercise, however. to show that (K*2). the postulate of SUCCESS, all by itself. implies that the theory $(K * A) \cap K$—which. note. is a consistent subtheory of K not containing the K-theorem $\neg A$—permits recovery of K via $\neg A$:[6]

> Let φ be an arbitrary theorem of K. Then K proves $\neg A \rightarrow \varphi$. By SUCCESS. $K * A$ proves A: whence $K * A$ proves $\neg A \rightarrow \varphi$. Hence $(K * A) \cap K$ proves $\neg A \rightarrow \varphi$. So $(K * A) \cap K. \neg A$ proves φ. But φ was an arbitrary theorem of K. Thus $(K * A) \cap K. \neg A$ permits recovery of K via $\neg A$.

Thus $K = [(K * A) \cap K. \neg A]$. So $(K * A) \cap K$ has all the properties that *AGM*-theory requires of a contraction of K with respect to $\neg A$. So it is reasonable to suspect that (K*6) does not do much to ensure minimal mutilation by revision.

2.2. The supplementary postulates. There are also two 'supplementary' revision postulates. (K*7) and (K*8):

(K*7) $K * (A \wedge B) \subseteq [K * A. B]$.
(K*8) If $K * A$ is consistent with B. then $[K * A. B] \subseteq K * (A \wedge B)$.

These supplementary postulates seek only to constrain the relationship between revision with respect to any conjunction. and revision with respect to either of its conjuncts.

It is well known that the supplementary postulates do not rule out the possibility of the 'amnesiac' revision. whereby $K * A$ is simply taken to be $[A]$.[7] So the supplementary postulates certainly do not speak to the problem of mutilation.

It may be suggested. however. that the supplementary postulates were intended. at least in part. to avert or at least to mitigate the unwelcome prospect of unnecessary *bloatings* by revision-functions.[8] But such a suggestion would be in error. The present paper shows that the supplementary postulates cannot fulfil such an

[6]This was first shown in Makinson [1987].

[7]This result is due to Mark Ryan [1996]. See footnote 9 for more relevant information in light of our generalization of Ryan's result by means of Theorem 3 below.

[8]Such a suggestion would flow naturally from the quotation above from p. 82 of Gärdenfors [1988]. See also Gärdenfors [1982]. at p. 142. where. in discussing the supplementary postulates for revision. he clearly sees those postulates as seeking to ensure that $K * (A \wedge B)$ is 'the *minimal* change of K necessary to include both A and B'. [Emphasis added.]

intention. As its title implies, it furnishes a crescendo of degeneracy results (Theorems 3, 4 and 5 below) for the full *AGM*-theory of belief-revision, supplementary postulates included. The result shows that the full set of postulates for the *AGM*-theory of rational belief-revision is so lax that revisions as irrational as can be imagined—both maximally mutilating and maximally bloating—can be delivered by revision-functions satisfying the full set.

§3. **Technical results.** The following definition is crucial to the results of this paper.

DEFINITION 7. *For C a sentence, and Θ any set of sentences*:

$$C \cdot \Theta =_{\text{df}} \begin{Bmatrix} [C, \Theta] & \textit{if } C, \Theta \nvdash \bot \\ [C] & \textit{if } C, \Theta \vdash \bot \end{Bmatrix}.$$

Our first aim is to prove the following central result (Theorem 3 below):

Let Θ be a consistent set of sentences. For any consistent theory K and any contingent sentence A refuted by K, set $K * A =_{\text{df}} A \cdot \Theta$. Then $*$ is a revision-function satisfying the AGM-*postulates* (K*1)–(K*8).

OBSERVATION 1. *$A \cdot \Theta$ is a theory.*

OBSERVATION 2. $A \cdot \Theta \vdash A$.

OBSERVATION 3. $A \cdot \emptyset = [A]$.

OBSERVATION 4. *If $A \nvdash \bot$, then $A \cdot \Theta \nvdash \bot$.*

OBSERVATION 5. *If $A \dashv\vdash B$ then $A \cdot \Theta = B \cdot \Theta$.*

OBSERVATION 6. $K \cap (A \cdot \Theta), \neg A \vdash K$.

PROOF. Let φ be an arbitrary K-theorem. Then $K \vdash \neg A \rightarrow \varphi$. By Observation 2, $A \cdot \Theta \vdash A$, whence $A \cdot \Theta \vdash \neg A \rightarrow \varphi$. Hence $K \cap (A \cdot \Theta) \vdash \neg A \rightarrow \varphi$. So $K \cap (A \cdot \Theta), \neg A \vdash \varphi$. But φ was an arbitrary K-theorem. Hence $K \cap (A \cdot \Theta), \neg A \vdash K$. ⊣

LEMMA 1. $\dfrac{A, \Theta \vdash \bot}{A \wedge B \cdot \Theta = [A \cdot \Theta, B]}$

PROOF. Suppose $A, \Theta \vdash \bot$. Then by $\wedge$L we have $A \wedge B, \Theta \vdash \bot$. By definition of $\cdot$, it follows that $A \wedge B \cdot \Theta = [A \wedge B]$. But $[A \wedge B] = [[A], B]$. Hence by transitivity of identity we have

(1) $A \wedge B \cdot \Theta = [[A], B]$.

We are supposing $A, \Theta \vdash \bot$. So by definition of $\cdot$ again, it follows that $A \cdot \Theta = [A]$. Substituting by means of this identity in (1) we obtain $A \wedge B \cdot \Theta = [A \cdot \Theta, B]$. ⊣

LEMMA 2. $\dfrac{A \wedge B, \Theta \nvdash \bot}{A \wedge B \cdot \Theta = [A \cdot \Theta, B]}$

PROOF. Suppose $A \wedge B, \Theta \nvdash \bot$. Then by $\wedge$L it follows that $A, \Theta \nvdash \bot$. Hence by definition of $\cdot$ we have $A \cdot \Theta = [A, \Theta]$. We are supposing $A \wedge B, \Theta \vdash \bot$, whence by definition of $\cdot$ again it follows that $A \wedge B \cdot \Theta = [A \wedge B, \Theta]$. But $[A \wedge B, \Theta] = [[A, \Theta], B]$. Hence $A \wedge B \cdot \Theta = [[A, \Theta], B]$. Substituting for $[A, \Theta]$ we obtain $A \wedge B \cdot \Theta = [A \cdot \Theta, B]$. ⊣

LEMMA 3. $\dfrac{A, \Theta \nvdash \bot \quad A \wedge B, \Theta \vdash \bot}{A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]}$

PROOF. Suppose both

(1) $A, \Theta \nvdash \bot$

and

(2) $A \wedge B, \Theta \vdash \bot$.

From (2) it follows by definition of $\cdot$ that

(3) $A \wedge B \cdot \Theta = [A \wedge B]$.

As a matter of logic we have $[A \wedge B] \subseteq [[A, \Theta], B]$. Substituting in this by means of identity (3) we obtain

(4) $A \wedge B \cdot \Theta \subseteq [[A, \Theta], B]$.

From (1) it follows by definition of $\cdot$ that $A \cdot \Theta = [A, \Theta]$. Substituting by means of this identity in (4) we obtain $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$. ⊣

LEMMA 4. $\dfrac{A, \Theta \nvdash \bot}{A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]}$

PROOF. Suppose $A, \Theta \nvdash \bot$. We proceed by constructive dilemma.

Case (i): $A \wedge B, \Theta \vdash \bot$. From this and our main supposition it follows by Lemma 3 that $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$.

Case (ii): $A \wedge B, \Theta \nvdash \bot$. From this it follows by Lemma 2 that

$$A \wedge B \cdot \Theta = [A \cdot \Theta, B],$$

whence obviously $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$. ⊣

THEOREM 1. $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$.

PROOF. We proceed by constructive dilemma.

Case (i): $A, \Theta \vdash \bot$. By Lemma 1 we have $A \wedge B \cdot \Theta = [A \cdot \Theta, B]$, whence obviously $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$.

Case (ii): $A, \Theta \nvdash \bot$. By Lemma 4 we have $A \wedge B \cdot \Theta \subseteq [A \cdot \Theta, B]$. ⊣

THEOREM 2. *If* $A \cdot \Theta, B \nvdash \bot$, *then* $A \wedge B \cdot \Theta = [A \cdot \Theta, B]$.

PROOF. Suppose $A \cdot \Theta, B \nvdash \bot$. We proceed by constructive dilemma.

Case (i): $A, \Theta \vdash \bot$. By Lemma 1, the result is immediate:

$$A \wedge B \cdot \Theta = [A \cdot \Theta, B].$$

Case (ii): $A, \Theta \nvdash \bot$. By definition of $\cdot$ it follows that $A \cdot \Theta = [A, \Theta]$. Substituting by means of this identity in our main supposition we obtain

(1) $[A, \Theta], B \nvdash \bot$.

Now, if it were the case that $A \wedge B, \Theta \vdash \bot$, we would have $[A, \Theta], B \vdash \bot$, contradicting (1). So $A \wedge B, \Theta \nvdash \bot$. It now follows by Lemma 2 that

$$A \wedge B \cdot \Theta = [A \cdot \Theta, B]. \quad ⊣$$

THEOREM 3. *Let* Θ *be a consistent set of sentences. For any consistent theory* K *and any contingent sentence* A *refuted by* K, *set*

$$K * A =_{\text{df}} A \cdot \Theta.$$

Then $*$ *is a revision-function satisfying the* AGM-*postulates* (K^*1)–(K^*8).

PROOF. By Observation 1, postulate (K*1) is satisfied: $K * A$ is a theory. By Observation 2, postulate (K*2) is satisfied: A is in $K * A$. (Postulate (K*3) becomes irrelevant, since we are concerning ourselves here only with the principal case for revision, in which K refutes A.) By Observation 4, postulate (K*4) is satisfied:

$K * A$ is consistent. By Observation 5. postulate (K*5) is satisfied: if A and B are logically equivalent. then $K * A = K * B$. Postulate (K*6) is satisfied (or at least satisfiable) since $K \cap (A \cdot \Theta)$ satisfies the postulates for being a contraction of K with respect to $\neg A$—in particular. by Observation 6. RECOVERY is satisfied. By Theorem 1. supplementary postulate (K*7) is satisfied: and by Theorem 2. supplementary postulate (K*8) is satisfied. ⊣

OBSERVATION 7. *With $\emptyset$ for Θ. Theorem 3 specializes to the result that the full set of postulates of the* AGM-*theory of revision-functions can be satisfied by always taking $K * A$ to be just the 'amnesiac' revision $[A]$ (which of course is maximally mutilating).*[9] *The definition of $K * A$ as $A \cdot \Theta$ (which is defined without reference to K) makes the deviant revision-function amnesiac about K in a similar way.*

Theorem 3 affords a uniform way of making potentially bizarre revisions. Note that Θ is an arbitrary consistent set of sentences. For many a theory K ripe for revision with respect to A. the theory $A \cdot \Theta$ will be bizarre as a revision of K with respect to A. This holds even (indeed. especially) when A is consistent with Θ. so that $A \cdot \Theta$ is $[A. \Theta]$. For then Θ will often effect unwanted bloating in addition to unwanted mutilation. This problem of bloating is particularly acute when Θ is a complete theory. or when Θ has no extra-logical vocabulary (that is. propositional atoms: or names. function signs and predicates) in common with K or with A.

If this were not already bad enough for the AGM-theory of revision. it is worth pointing out also that one can appeal to Theorem 3 so as to ensure that. for any given K ripe for revision with respect to A. and for any theory J that would definitely count as bizarre *qua* revision of K with respect to A. there will be an AGM-revision-function $*$ such that $K * A = J$. The way to ensure this will be described presently (see Theorem 4 below). after the following two Frobenian definitions.

DEFINITION 8. *Let D be a sentence and let H be a set of sentences. Then* $D \rightarrow H =_{\text{df}} \{D \rightarrow \varphi \mid \varphi \in H\}$.

Recall that by Definition 2. if J and K are sets of sentences then $J \vdash K$ holds just in case for every sentence $\varphi \in K$. $J \vdash \varphi$.

Fix a consistent theory H and a contingent sentence D such that $H \vdash D$.

LEMMA 5. $H \vdash D \rightarrow H$.

PROOF. Let φ be an arbitrary member of H. Then $H \vdash \varphi$. whence $H \vdash D \rightarrow \varphi$. So. by Definition 8 (of $D \rightarrow H$). we have $H \vdash D \rightarrow H$. ⊣

LEMMA 6. $D. D \rightarrow H \vdash H$.

PROOF. Let φ be an arbitrary member of H. By reflexivity $D \rightarrow \varphi \vdash D \rightarrow \varphi$. So by dilution on the left $D \rightarrow H \vdash D \rightarrow \varphi$. Hence $D. D \rightarrow H \vdash \varphi$. But φ was an arbitrary member of H. Hence $D. D \rightarrow H \vdash H$. ⊣

LEMMA 7. $H \vdash [D. D \rightarrow H]$.

[9] It was pointed out in Alchourrón and Makinson [1982] that 'full-meet revision' (as it subsequently came to be called) is this amnesiac revision operation. See their Observation 2.2 on p.19. The empty-set special case of Observation 7 of the present paper. however. is a stronger result. and was first shown by Ryan [1996]. See his Proposition 2. p. 132. Observation 7 generalizes this result of Ryan. with arbitary Θ in place of $\emptyset$. Both Ryan's result. and our generalization of it. take the supplementary postulates explicitly into account.

PROOF. Remember that our consistent H and contingent D have been fixed so that $H \vdash D$. By Lemma 5. we have $H \vdash D \rightarrow H$. Hence $H \vdash [D. D \rightarrow H]$. ⊣

LEMMA 8. $H = [D. D \rightarrow H]$.

PROOF. Immediate from Lemmas 6 and 7. ⊣

LEMMA 9. $[D. D \rightarrow H] \nvdash \bot$

PROOF. By Lemma 7. we have $H \vdash [D. D \rightarrow H]$. Suppose for *reductio* that $[D. D \rightarrow H] \vdash \bot$. Hence by multiple CUT we would have $H \vdash \bot$. contrary to H's assumed consistency. So $[D. D \rightarrow H] \nvdash \bot$. ⊣

LEMMA 10. $D \cdot D \rightarrow H = H$.

PROOF. By Lemma 9. we have $[D. D \rightarrow H] \nvdash \bot$. By definition of $\cdot$ it follows that $D \cdot D \rightarrow H \vdash [D. D \rightarrow H]$. By Lemma 8. $H = [D. D \rightarrow H]$. Hence $D \cdot D \rightarrow H = H$. ⊣

THEOREM 4. *Let J be a consistent theory. let D be a contingent sentence inconsistent with J. and let H be any theory that implies D. (H may be highly irrational as a 'revision' of J with respect to D: for H may involve both unwanted mutilation and unwanted bloating.) For any consistent theory K and any contingent sentence A refuted by K. set*

$$K * A =_{\mathrm{df}} A \cdot D \rightarrow H.$$

Then $*$ satisfies the AGM-*postulates* (K^*1)–(K^*8). *and* $J * D = H$.

PROOF. Let J. D and H be as specified in the hypotheses. By Theorem 3. the revision-function $*$ that we have defined satisfies the *AGM*-postulates (K^*1)–(K^*8). (Take $D \rightarrow H$ for Θ.) By Lemma 10. $K * D = H$. for all theories K ripe for revision with respect to D. Hence. in particular. $J * D = H$. ⊣

Theorem 4 tells us that a particular bizarre theory-revision can be extended to an overall revision-function satisfying the full set of *AGM*-postulates. This result can be strengthened (see Theorem 5 below). so as to provide infinitely many bizarre revisions for each theory K.

DEFINITION 9. *A* spectrum *consists of two lists of the same (finite or countably infinite) order-type $\gamma \leq \omega$:*

$$\{D_0. D_1. \dots\}$$
$$\{H_0. H_1. \dots\}$$

of contingent sentences D_i and consistent theories H_i. where for each $i < \gamma$. $H_i \vdash D_i$ and for $i \neq j < \gamma$. H_j refutes D_i.

EXAMPLE. For D_i take the sentence 'there are exactly i ψs' (for some consistent predicate ψ). and let H_i be any consistent theory containing D_i.

LEMMA 11. *For any spectrum $\langle\{D_i\}_{i<\gamma}.\{H_i\}_{i<\gamma}\rangle$. where $\gamma \leq \omega$. and for every $i < \gamma$. we have $H_i \vdash \cup_{j<\gamma}(D_j \rightarrow H_j)$.*

PROOF. Let a spectrum be given as described. Choose an arbitrary index $i < \gamma$. Let $j < \gamma$ be an arbitrary index. We proceed by constructive dilemma.

Case (i): $j = i$. By Lemma 5 we have $H_j \vdash D_j \rightarrow H_j$. Substituting j for i. we obtain $H_i \vdash D_j \rightarrow H_j$.

Case (ii): $j \neq i$. By the definition of a spectrum. we have $H_i. D_j \vdash \bot$. Hence $H_i \vdash D_j \rightarrow H_j$.

We now have $H_i \vdash D_j \rightarrow H_j$ for arbitrary $j < \gamma$. Hence $H_i \vdash \cup_{j<\gamma}(D_j \rightarrow H_j)$. ⊣

Lemma 12. *For any spectrum* $\langle\{D_i\}_{i<\gamma}, \{H_i\}_{i<\gamma}\rangle$, *where* $\gamma \leq \omega$, *and for every* $i < \gamma$, *we have* $H_i \vdash [D_i, \cup_{j<\gamma}(D_j \rightarrow H_j)]$.

Proof. By the properties of a spectrum, we have $H_i \vdash D_i$. By Lemma 11, we also have $H_i \vdash \cup_{j<\gamma}(D_j \rightarrow H_j)$. Hence $H_i \vdash [D_i, \cup_{j<\gamma}(D_j \rightarrow H_j)]$. ⊣

Lemma 13. *For any spectrum* $\langle\{D_i\}_{i<\gamma}, \{H_i\}_{i<\gamma}\rangle$, *where* $\gamma \leq \omega$, *and for every* $i < \gamma$, *we have* $D_i, \cup_{j<\gamma}(D_j \rightarrow H_j) \nvdash \bot$.

Proof. Let a spectrum be given as described. By Lemma 11 we have $H_i \vdash \cup_{j<\gamma}(D_j \rightarrow H_j)$. Suppose for *reductio* that $D_i, \cup_{j<\gamma}(D_j \rightarrow H_j) \vdash \bot$. Then by multiple cut we would have $H_i \vdash \bot$, contrary to the definition of a spectrum. Hence $D_i, \cup_{j<\gamma}(D_j \rightarrow H_j) \nvdash \bot$. ⊣

Lemma 14. *For any spectrum* $\langle\{D_i\}_{i<\gamma}, \{H_i\}_{i<\gamma}\rangle$, *where* $\gamma \leq \omega$, *and for every* $i < \gamma$, *we have* $D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j) \vdash H_i$.

Proof. Let a spectrum be given as described. By Lemma 13 we have

$$D_i, \cup_{j<\gamma}(D_j \rightarrow H_j) \nvdash \bot.$$

By definition of $\cdot$ it follows that

$$D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j) = [D_i, \cup_{j<\gamma}(D_j \rightarrow H_j)].$$

By Lemma 6 we have

$$D_i, D_i \rightarrow H_i \vdash H_i.$$

Hence by dilution on the left we have

$$[D_i, \cup_{j<\gamma}(D_j \rightarrow H_j)] \vdash H_i.$$

Substituting on the left of this by means of the last identity, we obtain

$$D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j) \vdash H_i.$$ ⊣

Lemma 15. *For any spectrum* $\langle\{D_i\}_{i<\gamma}, \{H_i\}_{i<\gamma}\rangle$, *where* $\gamma \leq \omega$, *and for every* $i < \gamma$, *we have* $H_i \vdash D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j)$.

Proof. Let a spectrum be given as described. By Lemma 13 we have

$$D_i, \cup_{j<\gamma}(D_j \rightarrow H_j) \nvdash \bot.$$

By definition of $\cdot$ it follows that

$$D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j) = [D_i, \cup_{j<\gamma}(D_j \rightarrow H_j)].$$

By Lemma 11 we have

$$H_i \vdash \cup_{j<\gamma}(D_j \rightarrow H_j).$$

Substituting on the right of this by means of the last identity we obtain

$$H_i \vdash D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j).$$ ⊣

Lemma 16. *For any spectrum* $\langle\{D_i\}_{i<\gamma}, \{H_i\}_{i<\gamma}\rangle$, *where* $\gamma \leq \omega$, *and for every* $i < \gamma$,

$$H_i = D_i \cdot \cup_{j<\gamma}(D_j \rightarrow H_j).$$

Proof. Immediate from Lemmas 14 and 15. ⊣

DEFINITION 10. *Given a consistent theory K, a K*-spectrum *is a spectrum* $\langle\{D_i^K\}_{i<\gamma}, \{H_i^K\}_{i<\gamma}\rangle$ *(for some $\gamma \leq \omega$), where, as indicated by the superscripts, the choice of the ith member of each list depends on K, and for every $i < \gamma$, K refutes D_i^K (so that K is ripe for revision with respect to D_i^K). Such a spectrum is* somewhere [resp., everywhere] bizarre *just in case for some* [*resp. every*] $i < \gamma$, *the theory H_i^K would count as a bizarre revision of K with respect to D_i^K.*

EXAMPLE of a K-spectrum. Let K contain $\neg\exists x\psi x$ (for some consistent predicate ψ). For D_i^K take the sentence 'there are exactly i ψs', and let H_i^K be any consistent theory containing D_i^K.

THEOREM 5. *For every consistent theory K, let* $\langle\{D_i^K\}_{i<\gamma}, \{H_i^K\}_{i<\gamma}\rangle$ *be a K-spectrum. Then there is an* AGM *revision-function $*$ satisfying all the* AGM-*postulates for revision, such that*:

$$\textit{For every } K, \textit{for every } i < \gamma,\ H_i^K = K * D_i^K.$$

This is the case even if one chooses K-spectra that are everywhere bizarre.

PROOF. For Θ in Theorem 3, take $\cup_{j<\gamma}(D_j^K \rightarrow H_j^K)$. So $K * A$ is defined to be $A \cdot \cup_{j<\gamma}(D_j^K \rightarrow H_j^K)$; whence by Lemma 16, we have $K * D_i^K = H_i^K$. ⊣

COROLLARY 1. *Suppose that K implies that there are at most k ψs. For every $n > k$, let D_n^K be the claim 'there are exactly n ψs'. Then for each $n > k$, one can choose a different bizarre theory J_n^K (subject only to the condition that J_n^K be consistent with D_n^K) and take $[J_n^K, D_n^K]$ as the revision of K with respect to D_n^K.*

PROOF. For every $n > k$, K refutes 'there are exactly n ψs' ($=D_n^K$), so K is ripe for revision with respect to the same. For each $n > k$, choose a different bizarre theory J_n^K (subject only to the condition that J_n^K be consistent with D_n^K). To obtain the desired K-spectrum $\langle\{D_i^K\}_{i<\omega}, \{H_i^K\}_{i<\omega}\rangle$, take for H_n^K the theory $[J_n^K, D_n^K]$. Theorem 5 tells us that we can take $[J_n^K, D_n^K]$ as the revision of K with respect to D_n^K. ⊣

COROLLARY 2. *Suppose that $K \vdash \forall x\psi x$. For every $n > 0$, let D_n^K be the claim 'there are exactly n non-ψs'. Then for each $n > 0$, one can choose a different bizarre theory J_n^K (subject only to the condition that J_n^K be consistent with D_n^K) and take $[J_n^K, D_n^K]$ as the revision of K with respect to D_n^K.*

PROOF. For every $n > 0$, K refutes 'there are exactly n non-ψs' ($=D_n^K$), so K is ripe for revision with respect to the same. The proof proceeds from here exactly as did the proof for Corollary 1. ⊣

With both these Corollaries, the bizarre extra beliefs J_n^K could, for example, differ in saying, among other things, that at any given time there are exactly n angels dancing on the head of a pin.

Corollary 2 shows that the full set of eight *AGM* postulates for revision admits of revision functions * such that for every theory K containing a universal generalzation, and for infinitely many A with respect to which K is ripe for revision, the revision $K * A$ can be as bizarre as one wishes.

At this juncture we reprise the point made at the end of §1.1. Note that Theorem 5 and its two Corollaries do not presuppose, or depend on, any particular analysis or formal explication of the notions of (minimal) mutilation or (minimal) bloating—or of bizarreness (*qua* revision). The dialectical structure of the predicament in which

AGM-theory is revealed to stand is as follows. The asserter of Theorem 5 is in effect saying '*You*, the listener, may choose whatever *you* would say is a bizarre series of revisions: then *I*, the speaker, will show you a revision function that satisfies all the *AGM* postulates, yet yields all those bizarre revisions.' The asserter of Theorem 5 is allowing his listener to work with any intuitive sense of bizarreness that she wishes. *She* gets to specify exactly what would count as a bizarre K-spectrum, *by her own lights*. (She might do this intuitively, without recourse to any particular logical or set-theoretic explication of the notions of minimal bloating and minimal mutilation.) Once she has a K-spectrum $\langle\{D_i^K\}_{i<\gamma}, \{H_i^K\}_{i<\gamma}\rangle$ for which she agrees that H_i would be bizarre as a revision of K with respect to D_i, Theorem 5 can immediately be applied so as to produce a revision function $*$ satisfying all eight AGM postulates, but such that $K * D_i = H_i$. Moreover, as Corollary 2 shows, any theory K containing a universally quantified claim $\forall x\psi x$ affords one the opportunity to create such spectral mischief. One need only contemplate revisions of K with respect to the infinitely various claims to the effect that there are, respectively, exactly $1, 2, 3\dots$ counterexamples to the universal generalization in question.

§4. **Concluding remarks.** In light of the foregoing results it would appear that an undesirable degree of laxity has been revealed in the full set of *AGM*-postulates for revision-functions. This laxity has been revealed by attention to the postulates themselves, rather than to the mathematical constructions afforded by the representation theorem.

4.1. Bizarre revisions as partial meet revisions. Naturally, because of the representation theorem about partial meet revisions, the question arises how one is to understand the possibility of these bizarre revisions in terms of partial meets. Daniel Osherson [2005] has provided a straightforward answer. The following account is for the case where one is concerned to satisfy only the basic postulates, and yet achieve a bizarre revision. (Note, however, that the account applies only to revisions that are closed under classical logic.)

THEOREM 6. *Suppose the theory K is ripe for revision with respect to A. Let J be any* classically closed *theory implying A. Then J is a partial meet revision of K with respect to A: that is, for some non-empty $\Gamma \subseteq (K \perp \neg A)$ we have $[\bigcap\Gamma, A] = J$.*

PROOF. Assume the hypotheses of the theorem. Consider $A \to J$, the set of conditionals of the form $A \to \varphi$, where $J \vdash \varphi$. K refutes A, hence proves any conditional of the form $A \to \varphi$. It follows that $(A \to J) \subseteq K$. For the sought Γ that the theorem claims to exist take the set of maximal non-($\neg A$-implying) subsets of K that include $A \to J$. Consider now $[\bigcap\Gamma, A]$. Clearly this is a partial meet revision of K with respect to A. We show that $[\bigcap\Gamma, A] = J$, by showing

(i) $\bigcap\Gamma, A \vdash J$—that is, for all C if $J \vdash C$, then $\bigcap\Gamma, A \vdash C$; and
(ii) $J \vdash [\bigcap\Gamma, A]$—that is, for all C if $\bigcap\Gamma, A \vdash C$, then $J \vdash C$.

Ad (i). By definition of Γ we have $A \to J \subseteq \bigcap\Gamma$. By Lemma 6 we have $A \to J, A \vdash J$. Hence by dilution we have $\bigcap\Gamma, A \vdash J$, as required.

Ad (ii). Suppose that $\bigcap\Gamma, A \vdash C$, where C is arbitrary. Then

$$\bigcap\Gamma \vdash A \to C \dots (1)$$

Since K refutes A, we have $K \vdash A \to \neg C$. Hence

$$[A \to J, A \to \neg C] \subseteq K.$$

Assume for *reductio* that

(2) $A \to J, A \to \neg C \vdash \neg A$; and
(3) $J \nvdash C$.

Since $J \vdash A \to J$, CUT on (2) yields $J, A \to \neg C \vdash \neg A$. Hence $J, \neg C \vdash \neg A$. But $J \vdash A$. So, since J is classically closed, we have $J \vdash C$, contradicting (3). We conclude

$$A \to J, A \to \neg C \nvdash \neg A,$$

and discharge our *reductio* assumption (2). Now expand $[A \to J, A \to \neg C]$ to a maximal non-($\neg A$-implying) subset γ of K. We have $\bigcap \Gamma \subseteq \gamma$. Hence by dilution on (1) we have $\gamma \vdash A \to C$. But we also have $\gamma \vdash A \to \neg C$. Hence $\gamma \vdash \neg A$, contradicting γ being non-($\neg A$-implying). We now conclude that $J \vdash C$, and discharge our *reductio* assumption (3). ⊣

The foregoing construction of a partial meet revision is perfectly general (in the classical case). So, even if J is judged to be bizarre as a revision of K with respect to A, we nevertheless know how to obtain J as a partial meet revision. Osherson shows also how to induce the necessary transitive relation on the power set of K so that a similar construction can be given of partial meet revisions that will satisfy the supplementary postulates as well.

4.2. Matters for further discussion. A fuller discussion will be undertaken elsewhere of possible interpretations of the results of this paper. Such a discussion needs to address various methodological and philosophical issues that we cannot enter into here. These issues concern the aims and methods involved when one seeks to provide a 'normative model' of phenomena such as logical inference, computability, and rational belief revision.

The question arises whether sensible further constraints on belief-revison functions can be laid down by further postulates within the *AGM* framework. A proposal in this direction has been made by Parikh [1999], who puts forward a 'splitting axiom' that might serve to curb or prevent bloating (see Axiom P, *loc. cit.*, p. 270). The axiom is stated here in the notation of this paper.

> Suppose K can be decomposed into $[K_1, K_2]$ in disjoint languages L_1, L_2 respectively. Suppose A is a formula of L_1. Then $K * A = [K_1 * A, K_2]$.

It remains to be investigated to what extent this axiom, and/or others like it, might mitigate the problem of degeneracy. But if the language of A is at all extensive, it is not clear that degeneracy can be avoided altogether. Even though Axiom P might help prevent one's bloating revisions from being utterly surreal, there is nevertheless no apparent reason to think that it will succeed in completely eliminating the possibility of revisions that are both mutilating and bloating (albeit only within the more restricted language L_1 in question).

REFERENCES

CARLOS ALCHOURRÓN, PETER GÄRDENFORS, AND DAVID MAKINSON [1985], *On the logic of theory change: partial meet contractions and revision functions*, this JOURNAL, vol. 50, pp. 510–530.

CARLOS ALCHOURRÓN AND DAVID MAKINSON [1982]. *On the logic of theory change: contraction functions and their associated revision functions*. ***Theoria***. vol. 48. pp. 14–37.

PETER GÄRDENFORS [1982]. *Epistemic importance and minimal changes of belief*. ***Australasian Journal of Philosophy***. vol. 62. pp. 136–157.

PETER GÄRDENFORS [1988]. ***Knowledge in flux***. MIT Press. Cambridge. MA.

DAVID MAKINSON [1987]. *On the status of the postulate of recovery in the logic of theory change*. ***Journal of Philosophical Logic***. vol. 16. pp. 383–394.

DANIEL OSHERSON [2005]. *Note on an observation by Neil Tennant*. Unpublished typescript, `http://www.princeton.edu/~osherson/papers/tennant11.pdf`.

ROHIT PARIKH [1999]. *Belief revision and language splitting*. ***Logic, language and computation*** (L. Moss. J. Ginzburg. and M. de Rijke. editors). CSLI Lecture Notes. Stanford. pp. 266–278.

MARK D. RYAN [1996]. *Belief revision and ordered theory presentations*. ***Logic, action, information: Essays on logic in philosophy and artificial intelligence*** (André Fuhrmann and Hans Rott. editors). Walter de Gruyter. Berlin and New York. pp. 129–151.

NEIL TENNANT [2005]. *Contracting intuitionistic theories*. ***Studia Logica***. vol. 80. pp. 371–393.

NEIL TENNANT [to appear]. *New foundations for a relational theory of theory-revision*. ***Journal of Philosophical Logic***.

DEPARTMENT OF PHILOSOPHY
THE OHIO STATE UNIVERSITY
COLUMBUS. OHIO 43210. USA
E-mail: tennant.9@osu.edu