Recovering From Tennant's Attack on Recovery

Erik J. Olsson[1]

I was introduced to the AGM theory of Peter Gärdenfors and his associates, Carlos Alchourrón and David Makinson, via Sten Lindström, who gave a series of lectures at the philosophy department in Uppsala on Peter's book *Knowledge in Flux* shortly after it appeared. At the time I was still an undergraduate student. The AGM theory and Peter's book in particular came to play a major role in my subsequent philosophical thinking. What I found particularly fascinating was the concept of epistemic entrenchment, and in my master thesis (i.e., my Swedish so called D-uppsats) I tried to develop Peter's suggestion that entrenchment be understood in terms of information-theoretic concepts. Actually, I sent the thesis to Peter, who called me some days later saying that liked it, but that I should change this and that, and that he had this EU project (DRUMS) that he thought I should participate in. Suddenly I was in the business... I have the AGM theory also to thank for providing the motivation for my first real, published paper (Hansson and Olsson, 1995). In the joint paper with Sven Ove Hansson we compared the theory of contraction laid out by AGM with Isaac Levi's theory put forward in his 1991 book.

This is not to say that I have always agreed with Peter; in my doctoral thesis on coherence (Olsson, 1997, see also Hansson and Olsson, forthcoming) I objected to his attempt to supplement the AGM theory with an epistemological interpretation (his coherence interpretation). Similar criticism was delivered in Olsson (1997b). Later I have

[1] The note is based on Olsson (1998). I would like to thank David Makinson for his extremely useful comments on an earlier version.

come to doubt the adequacy of the AGM theory as a theory of belief revision (rather than as, say, a theory of counterfactual reasoning). While the AGM model is impeccable from a logical perspective, I now think that the distance between the abstract theory and real flesh and blood belief revision is simply too great. The AGM theory, and much of the theorizing that has been inspired by it, tends to lead a life of its own in idealized Platonic heaven.

This note however is not concerned with legitimate criticism of the AGM theory, but with what I think is illegitimate criticism. In his paper "On having bad contractions, or: no room for recovery" (Tennant, 1997), Neil Tennant continues the attack on the AGM theory that he initiated in Tennant (1994). The target of the paper is, more specifically, the AGM postulate of recovery.

Before we look at Tennant's alleged counter arguments to recovery, let us for the sake of completeness rehearse the basics of the AGM framework. In the AGM theory the beliefs of an agent are represented as a logically closed set of sentences, often called a belief set or a theory. There are three main kinds of belief change according to this model: expansion (+), revision (*) and contraction (÷). The expansion of a set A by a sentence α is defined simply as $A+\alpha = Cn(A\cup\{\alpha\})$, i.e., as the set of logical consequences of $A\cup\{\alpha\}$. It is assumed that the revision of A by α can be constructed according to the identity $A*\alpha=(A\div\neg\alpha)+\alpha$ (the "Levi identity"). The kind of contraction favored in the AGM theory is partial meet contraction as defined by the identity $A\div\alpha=\cap\gamma(A\perp\alpha)$, where $A\perp\alpha$ is the set of inclusion-maximal subsets of A that do not imply α and $\gamma(A\perp\alpha)$ a non-empty subset of $A\perp\alpha$, unless the latter is empty, in which case $\gamma(A\perp\alpha)=\{A\}$. There are two special cases of partial meet contraction worthy of mention: maxichoice contraction (where, for all α, γ selects exactly one element of $A\perp\alpha$, in which case $A\div\alpha\in A\perp\alpha$), and full meet contraction (where, for all α, γ selects all elements of $A\perp\alpha$). The function γ is normally interpreted as representing the cognitive preferences of the agent in the sense that $\gamma(A\perp\alpha)$ is the set of the *best* maximally α-non-

implying subsets of A. The recommendation of partial meet contraction is, then, to take the "meet of the best" theories, thereby withholding judgment between these theories.

A somewhat different model is obtained by representing a belief state as a (typically) non-closed set A of sentences. The elements of such a "belief base" are intended to be fundamental beliefs (or axioms for the theory $Cn(A)$). In that kind of model, the logical closure $Cn(A)$ of the belief base A is changed only as a an effect of the change of A.

The AGM trio proved a series of representation theorems for their constructions (Alchourrón, Gärdenfors and Makinson, 1985). For example, an operation $\div$ is an operation of partial meet contraction if and only if it satisfies six so-called basic Gärdenfors postulates. One of these postulates is recovery saying that $A \subseteq (A \div \alpha) + \alpha$. Recovery is motivated by a principle of information economy ("minimum mutilation"): no information should get lost if we were to contract by a belief and then immediately expand by the same belief. Despite its initial appeal, recovery has been found questionable by several researchers who in other respects essentially subscribe to the AGM model. Perhaps most notably, Isaac Levi (1991) has argued that some information is of no value and, therefore, need not be retained, which may lead to violations of recovery. It is also well-known that recovery does not hold for contraction of belief bases (nor, of course, for theory contraction generated by base contraction). For related issues, see Fuhrmann (1991) and Hansson and Rott (1995). For a recent attempt to identify the contexts in which recovery is applicable, see Makinson (1997).

In his first counterexample to recovery (p. 251), Tennant asks us to consider the theory $Cn(\alpha)$. Since α implies $\beta \rightarrow \alpha$, it follows, according to Tennant, that $Cn(\alpha) \div \beta \rightarrow \alpha = Cn(\varnothing)$, and recovery is violated, or so it would seem. However, if contraction is performed on the theory level, it simply does not follow that $Cn(\alpha) \div \beta \rightarrow \alpha = Cn(\varnothing)$. For example, the sentence $\neg\beta \rightarrow \alpha$ is in $Cn(\alpha)$ and can still be retained in $Cn(\alpha) \div \beta \rightarrow \alpha$. If, on the other hand, contraction is performed on the base $\{\alpha\}$ for $Cn(\alpha)$, then we have indeed a counterexample to recovery, but only for (theory contraction generated by) contraction of belief bases. As noted above, the fact that base

contraction does not conform to recovery is common knowledge. Tennant's second purported counterexample rests on a similar ambiguity and at best shows what we already knew, namely, that recovery fails for base contraction.

Tennant continues (p. 252) by arguing that recovery fails to give expression to the requirement of minimum mutilation. In his view, there are contraction operations that satisfy recovery but nonetheless fail to adhere to the intuitive idea of minimum mutilation. The theory $\{\alpha\rightarrow\beta|\beta\in B\}$ is, it is claimed, "an awful mutilation" of $Cn(B)$, and yet it represents a contraction of $Cn(B)$ by α that satisfies recovery. The operation in question can indeed be expressed as a full meet contraction of α from $Cn(B)$. Thus the example does show that full meet contraction, although validating recovery, may go beyond minimal mutilation and may indeed yield unreasonably great mutilation. These points were, however, already emphasized by the AGM trio in their 1985 paper. Recovery thus does not imply minimal or even moderate mutilation; but it may well be, as AGM suggest, a necessary condition for each. At the same time, it may be said that the intuitive idea of minimal mutilation can hardly function as sole criterion for determining a contraction. For if taken literally and as the only criterion, it would clearly yield maxichoice contraction; but as is well known, there are strong formal arguments against the maxichoice operation, and intuitively, to choose to believe in just one maximal theory when there are several equally good such theories available is gratuitous.

Tennant's most surprising claim (pp. 252-254) is that one of the fundamental AGM proofs is erroneous. The criticized proof (from Alchourrón and Makinson, 1982) is intended to establish that recovery holds for maxichoice contraction of theories, and it is short enough to be repeated here. Suppose that $\alpha\notin Cn(\varnothing)$, $\alpha\in A$ and $\beta\in A$. We need to show that $\beta\in Cn((A\div\alpha)\cup\{\alpha\})$. Since Cn includes tautological implication, it will suffice to show that $\neg\alpha\vee\beta\in Cn(A\div\alpha)$. But since $\beta\in A$ and A is a theory, $\neg\alpha\vee\beta\in A$, so if $\neg\alpha\vee\beta\notin Cn(A\div\alpha)$ then

(i) $\neg\alpha\vee\beta\notin A\div\alpha$, so

(ii) $\alpha \in Cn((A \div \alpha) \cup \{\neg\alpha \vee \beta\})$.

Hence, $\alpha \in Cn((A \div \alpha) \cup \{\neg\alpha\})$, so $\alpha \in Cn(A \div \alpha)$, contradicting $A \div \alpha \in A \perp \alpha$.

In his criticism of this proof, Tennant maintains A) that neither of the hypothesis $\alpha \notin Cn(\varnothing)$ nor $\alpha \in A$ nor $\beta \in A$ is used when the proof is "[c]leaned up as much as possible", B) that the step from (i) to (ii) is left unmotivated, and C) that (ii) does not, in fact, follow from (i) given the assumptions of the proof. The source of the problems seems to be that Tennant fails to see that $A \div \alpha$ is supposed to be a maximal α-non-implying subset of A, an assumption used in the step from (i) to (ii). That $A \div \alpha$ is in fact a maximal α-non-implying subset of A follows immediately from the definition of maxichoice contraction, together with the assumption that $\alpha \notin Cn(\varnothing)$ (which is thus used) to ensure that $A \perp \alpha$ is not empty. Since $A \div \alpha$ is a maximal α-non-implying subset of A, we know that for every $\chi \in A \backslash A \div \alpha$ we have $\alpha \in Cn((A \div \alpha) \cup \{\chi\})$. Hence to conclude (ii) we need only check that $\neg\alpha \vee \beta \in A \backslash A \div \alpha$. That $\neg\alpha \vee \beta \notin A \div \alpha$ is given by (i). That $\neg\alpha \vee \beta \in A$ follows from the assumption that $\beta \in A$ (which is thus used) and the background assumption that A is closed under classical consequence. Thus (ii) holds and the proof is correct. What remains valid in Tennant's criticism is the minor point that the assumption $\alpha \in A$ is unnecessary.

According to Tennant the AGM theory "has completely lost sight of the fact that contraction need not be construed as a many-one relation between pairs of the form <theory, proposition> and theories", for "[w]hy *should* we expect 'the' contraction $A \div \alpha$ to be single-valued" (p. 258, notation adapted)? Tennant does not mention the work of Lindström and Rabinowicz, who studied AGM-style belief revision from a relational point of view in their 1991 paper.

Curiously, Tennant thinks that the coolness with which his 1994 paper was received by leading belief revisionists reflects their embarrassment over the fallacies that he

claims to have found in the AGM theory (Tennant, 1997, p. 256, footnote 9). Tennant does not consider the hypothesis that it may derive from his provocative style of writing. The present paper represents no improvement in this unfortunate respect. We are told, for instance, that the search for representation theorems has a "near-hypnotic hold" on the minds of the AGM theorists, who, suffering from "the functionality fetish", have not been "thinking seriously about the accuracy of the modeling". Instead, they "seek surrogates" for recovery in order to prove simple representation theorems, and so on (pp. 257-8).

In conclusion, where Tennant sees only uncritical acceptance of AGM, there is in fact diversity and open-mindedness with on-going research in many different directions. Tennant has, moreover, yet to produce a valid counter argument to recovery not already to be found in the standard literature on belief revision. His only original contribution to that debate, the argument to the effect that recovery does not hold even for AGM maxichoice contraction, is, as we have seen, defective.

References

Alchourrón, C. E., P. Gärdenfors, and D. Makinson, On the Logic of Theory Change: Partial Meet Contraction and Revision Functions, Journal of Symbolic Logic, 50 (1985): 510-530.

Alchourrón, C. E., and D. Makinson, On the Logic of Theory Change: Contraction Functions and their Associated Revision Functions, Theoria, 48 (1982): 14-37.

Fuhrmann, A., Theory Contraction Through Base Contraction, Journal of Philosophical Logic, 20 (1991): 175-203.

Hansson, S. O., and H. Rott, How Not to Change the Theory of Theory Change, British Journal for the Philosophy of Science, 46 (1995): 361-380.

Hansson, S. O., and E. J. Olsson, AGM Contraction and Levi Contraction: A Comparison, Notre Dame Journal of Formal Logic, 35 (1995): 103-119.

Hansson, S. O., and E. J. Olsson, Providing Foundations for Coherentism, to appear in Erkenntnis.

Gärdenfors, P., Knowledge in Flux: Modeling the Dynamics of Epistemic States, Cambridge, Mass., 1988.

Levi, I., The Fixation of Belief and Its Undoing, Cambridge University Press, Cambridge, Mass., 1991.

Lindström, S., and W. Rabinowicz, Epistemic Entrenchment with Incomparabilities and Relational Belief Revision, pp. 93-126 in A. Fuhrmann and M. Morreau (eds.), The Logic of Theory Change, Springer-Verlag, Berlin, 1991.

Makinson, D., On the Force of Some Apparent Counterexamples to Recovery, pp. 475-481 in E. G. Valdés et al. (eds.), Normative Systems in Legal and Moral Theory, Festschrift for Carlos E. Alchourrón and Eugenio Bulygin, Duncker und Homblot, Berlin, 1997.

Olsson, E. J., Coherence: Studies in Epistemology and Belief Revision, PhD thesis, Uppsala University, 1997a.

Olsson, E. J., A Coherence Interpretation of Semi-revision, in S. O. Hansson (ed.) Special Issue on Non-prioritized Belief Revision, Theoria, 1997.

Olsson, E. J., "Review of Tennant's 'On Having Bad Contractions, or: no Room for Recovery'", Zentralblatt für Mathematik, 1998.

Tennant, N., Changing the Theory of Theory Change: Towards a Computational Approach, British Journal for the Philosophy of Science, 45 (1994): 865-897.

Tennant, N., On Having Bad Contractions Or: No Room for Recovery, Journal of Applied Non-Classical Logic, 7 (1997): 241-266.