Note on an observation by Neil Tennant*

Daniel Osherson
Princeton University

July 13, 2005

Neil Tennant (Tennant, 2005) has offered an important observation about the AGM theory of belief revision (Gärdenfors, 1988). We attempt to restate and demonstrate his result in a slightly different way. Fix a formal language $\mathcal{L}$ that embeds sentential logic. Given $K \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$, $K \perp \varphi$ denotes the class of maximally consistent subsets of K that do not imply φ. That is, $A \in K \perp \varphi$ iff $A \subseteq K$, $A \not\models \varphi$, and there is no $B \subseteq K$ such that $B \supset A$ and $B \not\models \varphi$.

(1) PROPOSITION: Let deductively closed $K \subseteq \mathcal{L}$ and $\varphi \in \mathcal{L}$ be such that $K \models \varphi$. Let $T \subseteq \mathcal{L}$ be any satisfiable theory that implies $\neg\varphi$. Then for some $\Gamma \subseteq K \perp \varphi$, T is logically equivalent to $\bigcap \Gamma \cup \{\neg\varphi\}$.

Inasmuch as $\bigcap \Gamma \cup \{\neg\varphi\}$ is an AGM-style revision of K in face of contradicting $\neg\varphi$, the proposition seems to show that *anything goes* in the AGM framework. The latter framework can also be characterized by the axioms presented in Gärdenfors (1988); accepting them as conditions on belief revision is equivalent to using partial meet contraction as above.[1] So the AGM axioms are revealed by the proposition to be unrestrictive, which is what Tennant shows by considering them directly.

To prove (1), let deductively closed $K \subseteq \mathcal{L}$, $\varphi \in \mathcal{L}$, and satisfiable $T \subset \mathcal{L}$ be given. Suppose:

(2) (a) $K \models \varphi$

*Contact information: osherson@princeton.edu. Thanks to Neil Tennant for careful reading and comment.

[1]For discussion and proof, see Hansson 1999, pp. 65, 125ff.

(b) $T \models \neg\varphi$.

Let $A = \{\neg\varphi \rightarrow \psi \mid \psi \in T\}$. Then:

(3) (a) $T \models A$

(b) $A \subseteq K$

(c) $A \not\models \varphi$.

That $T \models A$ is evident from the truth-functional definition of $\rightarrow$. For (3)b, observe that $\varphi \models A$ (again by the definition of $\rightarrow$), hence by (2)a, $K \models A$, which implies (3)b by the deductive closure of K. For (3)c, suppose that $A \models \varphi$. Then by (3)a, $T \models \varphi$ which in the presence of (2)b contradicts the satisfiability assumed for T. Let $\Gamma = \{B \in K \perp \varphi \mid A \subseteq B\}$. Then $\Gamma \neq \emptyset$ by (3)b,c via the usual construction for creating maximally consistent subsets of K that do not imply φ. Thus, $A \subseteq \bigcap\Gamma$ hence by *modus ponens*, $\bigcap\Gamma \cup \{\neg\varphi\} \models T$.

To prove the converse, it suffices to show that $T \models \bigcap\Gamma$ since $T \models \neg\varphi$ by (2)b. It thus suffices to show that for all $\chi \in \mathcal{L}$, if $\bigcap\Gamma \models \chi$ then $T \models \chi$, equivalently, if $T \not\models \chi$ then $\bigcap\Gamma \not\models \chi$. So choose $\chi \in \mathcal{L}$ such that:

(4) $T \not\models \chi$

Since $\gamma \models \bigcap\Gamma$ for all $\gamma \in \Gamma$, to show that $\bigcap\Gamma \not\models \chi$, it is enough to exhibit $\gamma \subseteq \mathcal{L}$ such that:

(5) (a) $A \subseteq \gamma$

(b) $\gamma \in K \perp \varphi$

(c) $\gamma \not\models \chi$.

Since K is deductively closed, (2)a yields:

(6) $\chi \rightarrow \varphi \in K$.

Suppose for a contradiction that $A \cup \{\chi \rightarrow \varphi\} \models \varphi$. Then $A \cup \{\neg\varphi\} \models \neg(\chi \rightarrow \varphi)$. Therefore by (3)a and (2)b, $T \models \neg(\chi \rightarrow \varphi)$. Since $\neg(\chi \rightarrow \varphi) \models \chi$, it follows that $T \models \chi$, contradicting (4). Therefore:

(7) $A \cup \{\chi \rightarrow \varphi\} \not\models \varphi$.

By (3)b, (6) and (7), $A \cup \{\chi \rightarrow \varphi)\}$ can be extended to some $\gamma \in K \perp \varphi$. Since $\gamma \not\models \varphi$, $\gamma \not\models \chi$. Thus, γ satisfies (5). ■

A less arbitrary way to choose $\Gamma \subseteq K \perp \varphi$ is to define a transitive binary relation R over $pow(K)$ and $\Gamma_{R,\varphi} \subseteq K \perp \varphi$ such that:

$$\text{for all } \varphi \in \mathcal{L}, \Gamma_{R,\varphi} = \{\gamma \in K \perp \varphi \mid \forall \tau \in K \perp \varphi, \tau R \gamma\}.$$

Revision of K in the face of contradicting $\neg\varphi$ may then be identified with $\bigcap \Gamma_{R,\varphi} \cup \{\neg\varphi\}$. Intuitively, $\Gamma_{R,\varphi}$ contains the most attractive theories that survive confrontation with $\neg\varphi$ (thinking of R as expressing epistemic preference). Revision of this kind is characterized by axioms that extend those legitimating use of arbitrary $\Gamma \subseteq K \perp \varphi$.[2]

Reliance on a transitive relation R can be shown to restrict the pattern of contractions provoked by different data, but it seems not to have much effect on revision. Given $K, T \subseteq \mathcal{L}$, call $\varphi \in \mathcal{L}$ a K, T-*disagreement* just in case $K \models \varphi$ and $T \models \neg\varphi$.

(8) PROPOSITION: Let deductively closed $K \subseteq \mathcal{L}$ and satisfiable $T \subseteq \mathcal{L}$ be given. Then there is a transitive relation R over $pow(K)$ such that for all K, T-disagreements φ, T is logically equivalent to $\bigcap \Gamma_{R,\varphi} \cup \{\neg\varphi\}$.

Of course, Proposition (8) implies Proposition (1). For proof, let deductively closed $K \subseteq \mathcal{L}$ and satisfiable $T \subseteq \mathcal{L}$ be given. Define:

$$A = \{\varphi \vee \psi \mid \varphi \text{ is a } K, T\text{-disagreement and } T \models \psi\}.$$

$$\mathcal{A} = \{\gamma \in K \perp \varphi \mid A \subseteq \gamma \text{ and } \varphi \text{ is a } K, T\text{-disagreement}\}.$$

It is easy to verify:

(9) (a) $A \subseteq K$

(b) $T \models A$.

Choose R to be any transitive ordering of $pow(K)$ that renders exactly the members of $\mathcal{A}$ maximal. (There are many such orderings, including one that induces two equivalences classes, $\mathcal{A}$ and its complement, with $\mathcal{A}$ on top.)

[2]See Gärdenfors (1988) and Hansson (1999), Ch. 2.

Now let K, T-disagreement φ be given. For a contradiction, suppose $A \models \varphi$. Then $T \models \varphi$ by (9)b. But $T \models \neg\varphi$ because φ is a K, T-disagreement. This exhibits T as unsatisfiable, contradicting the hypothesis of (8). Therefore $A \not\models \varphi$ so by (9)a, A can be extended to some member of $K \perp \varphi$. Hence, $\mathcal{A} \cap (K \perp \varphi) \neq \emptyset$. Therefore, $\Gamma_{R,\varphi}$ of (8) is $\bigcap(\mathcal{A} \cap (K \perp \varphi))$. Since A is a subset of every member of $\mathcal{A} \cap (K \perp \varphi)$, and $\bigcap(\mathcal{A} \cap (K \perp \varphi)) \neq \emptyset$, $A \subseteq \bigcap(\mathcal{A} \cap (K \perp \varphi))$. It follows immediately that $\bigcap(\mathcal{A} \cap (K \perp \varphi)) \cup \{\neg\varphi\} \models T$.

For the converse, it suffices to prove $T \models \bigcap(\mathcal{A} \cap (K \perp \varphi))$ inasmuch as $T \models \neg\varphi$ (because φ is a K, T-disagreement). So it suffices to show that T implies every $\chi \in \mathcal{L}$ that $\bigcap(\mathcal{A} \cap (K \perp \varphi))$ implies, equivalently, that if $T \not\models \chi \in \mathcal{L}$ then $\bigcap(\mathcal{A} \cap (K \perp \varphi)) \not\models \chi$. So let $\chi \in \mathcal{L}$ be such that $T \not\models \chi$. For all $\gamma \subseteq \mathcal{L}$, if $A \subseteq \gamma \in K \perp \varphi$ then $\gamma \in \mathcal{A} \cap (K \perp \varphi)$, and $\gamma \models \bigcap(\mathcal{A} \cap (K \perp \varphi))$. It therefore suffices to exhibit $\gamma \subseteq \mathcal{L}$ such that

(10) (a) $A \subseteq \gamma$

(b) $\gamma \in K \perp \varphi$

(c) $\gamma \not\models \chi$.

Since K is deductively closed and implies every K, T-disagreement, we have:

(11) $\chi \rightarrow \varphi \in K$.

Suppose for a contradiction that $A \cup \{\chi \rightarrow \varphi\} \models \varphi$. Then $A \cup \{\neg\varphi\} \models \neg(\chi \rightarrow \varphi)$. Hence, $T \models \neg(\chi \rightarrow \varphi)$ by (9)b and the fact that $T \models \neg\varphi$ (because φ is a K, T-disagreement). But $\neg(\chi \rightarrow \varphi) \models \chi$, so $T \models \chi$, contradicting the choice of χ. Hence,

(12) $A \cup \{\chi \rightarrow \varphi\} \not\models \varphi$.

From (9)a, (11), and (12), $A \cup \{\chi \rightarrow \varphi\}$ can be extended to some $\gamma \in K \perp \varphi$. Since $\gamma \not\models \varphi$, $\gamma \not\models \chi$. Thus, γ satisfies (10). ■

References

GÄRDENFORS, P. (1988): *Knowledge in Flux: Modeling the dynamics of Epistemic States*. MIT Press, Cambridge MA.

HANSSON, S. O. (1999): *A Textbook of Belief Updating*. Kluwer, Dordrecht.

TENNANT, N. (2005): "On the Degeneracy of the Full AGM-Theory of Theory-Revision," Discussion paper, Ohio State University.