Ranking Theory and Stalnaker's Counterexample

Nathan Bice

December 9, 2009

Abstract

In a recent article (Stalnaker, 2009), Robert Stalnaker presented a general counterexample to several popular accounts of iterated belief revision, including Ranking Theory. In this paper, I will argue that Stalnaker's counterexample fails as a counterexample to the standard Ranking-Theoretic account of iterated revision. In particular, I will argue that his counterexample cannot be expressed in a Ranking-Theoretic framework, and the most plausible modifications to his counterexample that allow it to be expressed in such a framework either do not result in the counterintuitive conclusion as Stalnaker claims or do result in this conclusion, but it is no longer counterintuitive. I will begin this paper with a brief introduction to AGM Belief Revision, followed by an introduction to Ranking Theory as motivated by problems in the AGM framework. I will then discuss a recent paper (Hild and Spohn, 2008) which attempts to give an operational definition of ranks in terms of iterated contraction functions via measurement theory. I will end with discussion of the Stalnaker counterexample.

Contents

1 A Brief Introduction to AGM Belief Revision

AGM Belief Revision gives us an account of rational constraints on changes in non-quantitative states of belief, states which include belief, disbelief, and suspension of judgment.[1] We'll begin our discussion by assuming that an agent's *belief set* $\mathcal{B}$ can be represented by a set of propositions in some Boolean Algebra $\mathcal{A}$ of subsets of a nonempty set of possibilities W.[2] So, beliefs are represented by their *contents*, contents are *propositions*, and propositions are represented by the set of possibilities in which they are true. These are all controversial assumptions, but they are not the subject of this paper. Throughout this paper, I will unabashedly call elements of $\mathcal{A}$ *propositions.*

Belief sets will be formally defined as follows: Let $\mathcal{A}$ be a Boolean Algebra over W. Then $\mathcal{B}$ is a *belief set* iff $\mathcal{B}$ is a subset of $\mathcal{A}$ such that $W \in \mathcal{B}$; $\emptyset \notin \mathcal{B}$; if $\mathcal{B}' \subseteq \mathcal{B}$ then $\cap\mathcal{B}' \in \mathcal{B}$; and if $A \in \mathcal{B}$ and $A \subseteq B$, then $B \in \mathcal{B}$. So, $\mathcal{B}$ contains the tautology, does not contain the contradiction, and is closed under both arbitrary intersection (conjunction) and the superset relation (logical consequence). For simplicity, we're assuming that belief sets are closed under *arbitrary* rather than finite intersection. That way, we can refer to what Wolfgang Spohn (unpublished) calls the *core* C of a belief set, where for any $A \in \mathcal{A}$, $A \in \mathcal{B}$ exactly if $C \subseteq A$. C is simply the set of all possibilities that are consistent with what the agent believes.

Let's proceed to the dynamics of belief. Now, in the case where the agent receives new information that is *consistent* with what the agent believes (in our propositional framework, the agent receives information (with propositional content) A such that $C \cap A$ is nonempty, where C is the agent's prior core), there is a standard account of what the

[1]As is standard, I will treat disbelief in a proposition as equivalent to belief in its negation. I assume that the reader has some familiarity with the standard AGM model of Belief Revision, as presented in (Gärdenfors, 1988). This section will only contain a brief summary of AGM Belief Revision, at the level of propositions rather than sentences and within a Ranking-Theoretic framework, following Wolfgang Spohn (unpublished). The purpose of this section is to provide motivation for Ranking Theory.

[2]W can be, for example, the set of outcomes of a single toss of a six-sided die or even the set of all possible worlds. Throughout, I will implicitly assume that certain subsets of W are in fact elements of the algebra $\mathcal{A}$. For an explicit account of the needed constraints on $\mathcal{A}$, see (Spohn, unpublished).

result of the agent's updating her beliefs ought to be; namely, the agent's posterior belief core C' is the intersection of the agent's prior belief core C and the (propositional content of the) new information A. Thus, after updating, the agent *believes* that D exactly if $C \cap A$ is a subset of D. So, after updating, the agent's posterior belief set $\mathcal{B}'$ consists of all of her prior beliefs, the new belief that A, and all beliefs resulting from closure under conjunction and logical consequence after A is added to her prior belief set $\mathcal{B}$. This is known as *expansion* in the AGM framework.[3]

But what if the agent receives new information A that is *inconsistent* with what the agent believes? Here are two obvious constraints that one could put on updating an agent's belief set in light of such information: The posterior core C' should be nonempty (after updating, there should be some possibility that is consistent with what the agent believes) and C' should be a subset of A (A should be believed after updating). Furthermore, one might consider generalizing our constraints on expansion in the following way: if, *after* updating on information A, some proposition B is consistent with what the agent *now* believes ($C' \cap B \neq \emptyset$), then if the agent *had* updated on $A \wedge B$ instead of simply A, the resulting core C'' should have been exactly $C' \cap B$, regardless of whether or not A *itself* was consistent with what the agent previously believed. Here is the motivation: If B is in fact consistent with the agent's beliefs after updating on A alone, then updating on A and then B should be the same as updating on $A \wedge B$ all at once. Furthermore, given our constraints on expansion, if B is consistent with C', then updating on B (after having already updated on A) should result in a posterior core $C'' = C' \cap B$. If we combine these constraints, we get what I will call a *Propostional AGM Revision Function* $*$, which is defined relative to an agent's prior belief set and maps propositions (representing the new information) to posterior belief cores. In particular, the agent's prior belief core will be identical to $*(W)$, the result of "learning" the tautological proposition W (i.e. vacuous expansion).[4]

[3] C is of course contracted, since by *expanding* her beliefs the agent *contracts* the set of possibilities consistent with what she believes.

[4] Spohn (1988) originally called such functions *simple selection functions.* Following Spohn, I have made two simplifying assumptions above that distinguish such functions from standard AGM revision

Definition. Let $\mathcal{A}$ be a Boolean Algebra over W. Then $*$ is a *Propositional AGM Revision Function* exactly if $*$ is a function from $\mathcal{A} - \{\emptyset\}$ into $\mathcal{A}$ such that for all $A, B \in \mathcal{A} - \{\emptyset\}$:

(a) $\emptyset \neq *(A) \subseteq A$,

(b) if $*(A) \cap B \neq \emptyset$, then $*(A \cap B) = *(A) \cap B$.

Now we can state what Spohn calls the *law of simple conditionalization*: if the propostional AGM revision function $*$ characterizes the doxastic state of the agent s at time t and if A is the (propositional content of the) total information s receives and accepts between t and t', then $*(A)$ is $s's$ belief core at t', so that s *believes* that D at t' exactly if $*(A) \subseteq D$.

There is an equivalent way of representing revision functions that will help elucidate the amount of structure these constraints impose on the agent's doxastic state at some time t. A revision function $*$ corresponds precisely to a transitive and complete weak well-ordering (WWO) $\preceq$ of the set of possibilities W. Importantly, although every non-empty proposition A of $\mathcal{A}$ has a minimum with respect to $\preceq$, since $\preceq$ is not necessarily antisymmetric, there may be more than one $\preceq$-minimal possibility in A. Now, given a $WWO \preceq$, one can construct a unique revision function $*$ by setting $*(A)$ equal to the set of all $\preceq$-minimal possibilities in A for each $A \in \mathcal{A}$. Similarly, given a revision function $*$, one can construct a unique $WWO \preceq$ by letting $*(W) = C_0$ be the $\preceq$-minimal possibilities in W, $*(W - C_0) = C_1$ be the $\preceq$-minimal possibilities in $(W - \min_{\preceq}(W)) = W_1$ (i.e., the second-most minimal possibilities in W), $*(W - (C_0 \cup C_1)) = C_2$ be the $\preceq$-minimal possibilities in $(W - (\min_{\preceq}(W) \cup \min_{\preceq}(W_1))) = W_2$ (i.e., the third-most minimal possibilities in W), etc.[5] One can also use these definitions to recover the original $\preceq$ or $*$.

functions: all belief sets are consistent, and agents cannot update on information which is *itself* contradictory. More importantly, AGM revision functions are standardly defined as functions from *any* belief set and information A to a posterior belief set, while I have followed Spohn in defining a revision function relative to a *particular* belief set. In this way, one can think of a revision function as characterizing the doxastic state of the agent at some particular time t without assuming that the agent will continue to revise her beliefs (or has always revised her beliefs) according to this revision function. In fact, we will soon see that there are reasons for rejecting these assumptions.

[5] I'm ommitting the full definition via transfinite recursion. See (Spohn, 1988) or (Spohn, unpublished)

Intuitively, a revision function $*$ corresponds to an *ordering of disbelief* $\preceq$ over the possibilities in W, where the agent's current belief core $C_0 = *(W)$ consists of the possibilities not disbelieved at all, $C_1 = *(W - C_0)$ (that is, the posterior core that would result from the agent updating by the information that some of the her beliefs are false) consists of the possibilities disbelieved the least, etc. So, when the agent receives new information A that is inconsistent with what she believes, she revises her core to the set of least-disbelieved possibilities that are consistent with A, that is, the $\preceq$-minimal possibilities in A.[6]

I will now define a *Propositional AGM Contraction Function* $\div$ before describing a serious shortcoming of the AGM framework, since contraction functions will become very important later on. Unfortunately, contraction functions don't fit as nicely into the story as I have been telling it (in terms of updating beliefs in light of new information), but intuitively they correspond to an agent's suspending judgment regarding a formerly-believed proposition A, so that after contraction by A, the agent neither believes A nor believes $\bar{A}$. The agent's prior belief core C will be equal to $\div(\emptyset)$; that is, C is equal to the posterior core after the vacuous contraction of the contradictory proposition (which wasn't believed in the first place). Formally, these functions are defined as follows:

Definition. Let $\mathcal{A}$ be a Boolean Algebra over W. Then $\div$ is a *Propositional AGM Contraction Function* exactly if $\div$ is a function from $\mathcal{A} - \{W\}$ into $\mathcal{A}$ such that for all $A, B \in \mathcal{A} - \{W\}$:

(a) $\emptyset \neq \div(\emptyset) \subseteq \div(A) \subseteq \div(\emptyset) \cup \bar{A}$ and $\div(A) \nsubseteq A$,

(b) if $\div(A \cap B) \nsubseteq A$, then $\div(A) \subseteq \div(A \cap B) \subseteq \div(A) \cup \div(B)$.

Condition (a) captures the intuition that, in order for the agent to give up her belief that A, she must add $\bar{A}$-possibilities to the set of possibilities consistent with what she

for details. This result is of course similar to the correspondence between standard AGM revision functions and *entrenchment orderings* as defined in (Gärdenfors, 1988).

[6] In (Gärdenfors, 1988), an entrenchment ordering is *not* considered an ordering of disbelief, since all elements of the agent's belief set are maximally believed, but rather an ordering of "usefulness in inquiry and deliberation". Such differences, while very important, are not relevant for our purposes.

believes, but she should not add any further A-possibilities (she shouldn't give up belief that D if D doesn't even entail A, unless $D \wedge \bar{A}$ is no less implausible than $\bar{D} \wedge \bar{A}$, and hence the D-possibilities were only less implausible than the $\bar{D}$-possibilities with respect to A), nor should she remove possibilities (she shouldn't gain *new* beliefs). Condition (b) captures the intuition that, if after contracting by $A \wedge B$ the agent does not believe that A, then contracting by A alone would have resulted in giving up no more beliefs than contraction by $A \wedge B$, which itself resulted in giving up no more beliefs than if the agent had kept only the beliefs that would remain no matter whether she had contracted by A alone or contracted by B alone. Contracting by B alone, however, may have resulted in giving up more beliefs than contracting by $A \wedge B$, since B was not necessarily contracted from the agent's belief set when she gave up her belief in $A \wedge B$.

Intuitively, when forced to contract her belief that A, the agent adds the least-disbelieved $\bar{A}$-possibilities to her core. Now, revision and contraction are interdefinable via the following theorem:[7]

Theorem 1. Let $\mathcal{A}$ be a Boolean Algebra over W. If $*$ is a revision function for $\mathcal{A}$, then the function $\div$ defined by $\div(A) = *(W) \cup *(\bar{A})$ is a contraction function for $\mathcal{A}$, and one can recover $*$ from $\div$ via the definition $*(A) = \div(\bar{A}) \cap A$. Furthermore, if $\div$ is a contraction function for $\mathcal{A}$, then the function $*$ defined by $*(A) = \div(\bar{A}) \cap A$ is a revision function for $\mathcal{A}$, and one can also recover $\div$ from $*$ via the previous definition.

[7] This theorem can be found in (Spohn, unpublished), although it is merely the propositional variant of a similar theorem in (Gärdenfors, 1988). The two identities are originally due to William Harper (1977) and Isaac Levi (1977), respectively. Now, if someone rejected the Recovery Postulate (in our framework, this is: if $\div(\emptyset) \subseteq A$, then $\div(A) \cap A \subseteq \div(\emptyset)$, which follows from condition (a)), perhaps because he didn't think one should always refrain from adding A-possibilities to the posterior core when contracting by A, he would no longer be compelled to consider contraction reducible to revision, although he may still be happy to accept that revision is reducible to contraction and expansion (Recovery is unnecessary for that half of the proof). Makinson (1987) has shown that, even if we reject Recovery, there will be a collection of functions that satisfy the other axioms for contraction such that, via the Levi Identity, they all result in the same revision function. Furthermore, there will be a *unique* such function that also satisfies Recovery, namely, the function in the collection that contracts the least amount of beliefs from the agent's belief set. I will not defend Recovery here, but instead refer the reader to (Spohn, unpublished).

2 Trouble in Paradise

This can't be the end of the story, however. There is a serious problem with the law of simple conditionalization as presented above: the prior doxastic state of the agent is represented by a revision function, while the posterior doxastic state of the agent is merely represented by a belief set. So far, we have given no account of how the agent revises her beliefs more than once. To put the point another way, we have seen how, given an ordering of disbelief, the agent revises her belief set to a new belief set in light of new information, but we haven't said anything about how the agent revises the ordering of disbelief *itself*! This is a variant of what is known in the literature as the *Problem of Iterated Belief Change*.

Couldn't we have defined revision functions more generally as applying to *all* belief sets and input propositions at once, that is, as a function from both a belief set and a proposition to another belief set, thus allowing the agent to continue revising her beliefs using the *same* revision function with respect to her posterior belief set? Certainly, but orderings of disbelief (and entrenchment orderings) correspond to *particular* belief sets, and hence we have still placed basically no constraints on how the agent moves from one ordering of disbelief to another ordering of disbelief. So far, we only have the trivial constraints that the posterior ordering must be a transitive and complete weak well-ordering and the agent's posterior core must be precisely the minimal elements of this new ordering, but this amounts to saying no more than the agent should still be updating via a revision function and the agent's posterior belief set should in fact be the agent's belief set when the agent updates again. At this point, almost anything goes.[8]

Do we have reasons for thinking that there *ought* to be additional constraints on iterated belief revision? Adnan Darwiche and Judea Pearl (1997) give several examples of

[8]In fact, defining revision functions in this general way might be too *restrictive* as well, since the agent would then be committed to *always* revising a particular belief set in a particular way. It isn't obvious that whenever the agent revises her beliefs and ends up with the same belief set she had at some previous time, her *dispositions* to revise this belief set in light of any new information should now be identical to what they were previously.

AGM-compatible revision functions that license highly counterintuitive revision behavior precisely because these functions allow for counterintuitive iterated revisions. I reproduce one such example here:

"*Example 1* We see a strange new animal X at a distance, and it appears to be barking like a dog, so we conclude that X is not a bird, and that X does not fly. Still, in the event that X turns out to be a bird, we are prepared to change our mind and conclude that X flies. Observing the animal closely, we realize that it actually can fly. The question now is whether we should retain our willingness to believe that X flies in case X turns out to be a bird after all. We submit that it would be strange to give up this conditional belief merely because we happened to observe that X can fly. Yet, we provide later an AGM-compatible revision operator $\circ$ that permits such behavior."

Now, let's consider a couple of obvious attempts to determine exactly how the agent's ordering of disbelief should evolve when she revises by information A that was previously inconsistent with what the agent believed, both mentioned in (Spohn, 1988). Perhaps all of the A-possibilities should now be shifted in the ordering so that they are all before all of the $\bar{A}$-possibilities, with the ordering otherwise preserved. However, since $\bar{A}$ was previously believed, this amounts to shifting *all* of the possibilities consistent with what the agent previously believed so that they are now at least as disbelieved as the previously most-disbelieved A-possibility. In many cases, this seems very counterintuitive. As Spohn (1988) puts it, this amounts to accepting information A with *maximum* firmness and is not suitable as a general model.

Perhaps only the $\preceq$-minimal A-possibilities should be shifted to the beginning of the ordering (as is forced by the revision axioms), while the rest of the ordering should remain unchanged. However, this amounts to accepting A with *minimum* firmness, since if the agent merely found out that at least one of her beliefs is false, she would immediately give up her belief that A and return to her previous belief set. Additionally, neither this proposal nor the previous proposal are *reversible* in the sense that there is no general way of inferring the previous ordering of disbelief from the new ordering of disbelief, even if we know what position the minimal A-possibilities occupied in the previous ordering.

I'll briefly explain why the first proposal is not reversible, since the reason for this is very illustrative of what the Ranking Theorist sees as a general problem with the account given so far (and with many other attempts in the literature to fix it). We can think of the ordering $\preceq$ as a well-ordered partition of the possibilities in W, where each element of the partition is a set of possibilities that are all in the same position in the ordering. The partition can be written as follows: $C_0, C_1, C_2, \ldots, C_n, \ldots$, where C_0 is the agent's core, C_1 is the set of possibilities disbelieved the least, C_2 is the set of possibilities disbelieved the second-least, etc. Now, assume that the least-disbelieved A-possibilities are in C_n and (for simplicity) there exist most-disbelieved A-possibilities that are in C_{n+m}. So, the first proposal amounts to shifting to the following partition: $C_n \cap A, \ldots, C_{n+m} \cap A, C_0, C_1, \ldots, C_{n-1}, C_n \cap \bar{A}, \ldots, C_{n+m} \cap \bar{A}, C_{n+m+1}, \ldots$. However, empty terms must be removed in order for this partition to correspond to an ordering. For example, if there are no A-possibilities in C_{n+3}, then $C_{n+3} \cap A$ will be empty and its position as the set of third-least disbelieved possibilities will instead be occupied by $C_{n+4} \cap A$. So, after the agent updates on new information A, even if we know the value of n, there is no general way of inferring exactly what the agent's previous ordering of disbelief was, since some information about relative positions in the ordering is lost.

Wolfgang Spohn (1988) claims that these examples demonstrate what must be done in order to make progress on the problem of iterated belief change: we must specify the *firmness* with which the agent incorporated new information into the new state, and we must not delete empty terms from the agent's well-ordered partition.[9] Hence, rather than merely characterizing the agent's doxastic state at some time t as an *ordering* of disbelief, we should characterize it as a *ranking* of disbelief, where certain *arithmetical* properties

[9]Why should we care about reversibility? Spohn (unpublished) defends reversibility as follows: if Wolfgang asserts E and shortly thereafter apologizes and takes E back, W may simply not be rich enough to appropriately express the content of the beliefs that "Wolfgang asserts E" and "Wolfgang says, "Sorry, I was wrong about E"" as propositions in the algebra. Hence, if W is not rich enough, the second revision is perhaps best thought of as undoing the first revision. For similar reasons, it may be best to consider the first revision as a revision by E with some firmness n, where n corresponds to something like Wolfgang's perceived trustworthiness, rather than a revision by "Wolfgang asserts E".

of the ranks are meaningful. In particular, if some possibility is disbelieved with rank 3 and another possibility is disbelieved with rank 5, then the difference of 2 matters, *even if* there are no possibilities disbelieved with rank 4. With this motivation in mind, we finally turn to Ranking Theory.

3 Ranking Theory

A *ranking function* assigns a natural number or ∞ to every possibility in W, where the natural numbers are meant to represent the agent's rank of disbelief in that possibility. So, for each n, $\kappa^{-1}(n)$ corresponds to the (possibly empty) set of possibilities disbelieved with rank n. More formally, κ is a *ranking function* for $\mathcal{A}$ iff κ is a function from W into $\mathbb{N}^+ = \mathbb{N} \cup \{\infty\}$ such that $\kappa^{-1}(0) \neq \emptyset$. Now, we extend ranks to *propositions* by defining $\kappa(\emptyset) = \infty$ and $\kappa(A) = \min\{\kappa(w)|w \in A\}$ for each non-empty $A \in \mathcal{A}$.

$\kappa^{-1}(0)$ is of course the agent's core, and hence requiring that this set be nonempty amounts to requiring that there is at least one possibility consistent with what the agent believes. The definition of the rank of a *proposition* is motivated by the idea that the agent's degree of disbelief in a proposition A should be precisely her degree of disbelief in what she considers to be the least implausible possibilities in which A is true. Now, the agent *believes* that A exactly if the ranking function characterizing the agent's doxastic state is such that $\kappa(\bar{A}) > 0$. If $\kappa(A) = \kappa(\bar{A}) = 0$ then the agent suspends judgment regarding A. One can then define the *belief set associated with a negative ranking function* κ, $\mathcal{B}(\kappa)$, as the set of all propositions $A \in \mathcal{A}$ such that $\kappa(\bar{A}) > 0$.

One can easily show that for any $A \in \mathcal{A}$, either $\kappa(A) = 0$ or $\kappa(\bar{A}) = 0$ or both. The motivation should be clear: either the agent does not disbelieve that A or the agent does not disbelieve that $\bar{A}$ or both, for otherwise the agent would both believe that A and believe that $\bar{A}$. Also, for any $A, B \in \mathcal{A}$, $\kappa(A \cup B) = \min\{\kappa(A), \kappa(B)\}$. In fact, this result holds for *arbitrary* (rather than merely *finite*) union.

One of the central definitions in Ranking Theory is the definition of the *conditional rank*: the rank of B under the supposition that A. Here are two considerations that motivate the definition: $\bar{A}$ should be maximally disbelieved under the supposition that A ($\kappa(\bar{A}|A) = \infty$), and the ordering of possibilities *within* A should not change under the supposition that A. So, let κ be a ranking function for $\mathcal{A}$ and $A \in \mathcal{A}$ such that $\kappa(A) < \infty$. Then for any $w \in W$ the *conditional rank of w given A* is defined as

$$\kappa(w|A) = \left\{ \begin{array}{ll} \kappa(w) - \kappa(A), & \text{if } w \in A \\ \infty, & \text{if } w \in \bar{A} \end{array} \right\}.$$

We then extend this definition to propositions in the obvious way, namely, by defining $\kappa(B|A) = \min\{\kappa(w|A)|w \in B\}$ when B is nonempty and $\kappa(\emptyset|A) = \infty$. As a consequence of this definition, if $\kappa(A) < \infty$, then $\kappa(B|A) = \kappa(A \cap B) - \kappa(A)$.

Now we can define the *conditionalization* of a ranking function κ by some proposition A with firmness $n \in \mathbb{N}^+$ (where $\kappa(A), \kappa(\bar{A}) < \infty$) as follows:

$$\kappa_{A \to n}(w) = \left\{ \begin{array}{ll} \kappa(w|A), & \text{if } w \in A \\ \kappa(w|\bar{A}) + n, & \text{if } w \in \bar{A} \end{array} \right\}.$$

The idea is that an agent incorporates new information into her doxastic state with varying degrees of firmness (thus the parameter n) and, after updating on A, the agent's *ordering* of disbelief *within* A and *within* $\bar{A}$ should remain the same, even though the *ranks* of disbelief will change and the ordering of A-possibilities and $\bar{A}$-possibilities will shift relative to each other.[10] For example, if the agent believes that $\bar{A}$ and then conditionalizes on A with firmness $n > 0$, all the A-possibilities will shift downward so that the minimal A-possibilities are now ranked 0 and in general the A-possibilities will shift from $\kappa(w)$ to $\kappa(w) - \kappa(A)$. Furthermore, in this case the $\bar{A}$-possibilities will shift upward so that $\kappa_{A \to n}(\bar{A}) = n$ and in general the $\bar{A}$-possibilities will shift from $\kappa(w)$ to $\kappa(w) + n$. So, after conditionalization, the agent will believe that A and will disbelieve that $\bar{A}$ with rank n.

Now we can state a *dynamic law of conditionalization* for ranking functions: If the prior doxastic state of the subject s at time t is characterized by the ranking function κ

[10]Darwiche and Pearl (1997) also recommended that the ordering within A and within $\bar{A}$ should not change when revising by A, and this seems to be generally accepted in the AGM literature on the problem of iterated belief change.

and if s receives and accepts information (with propositional content) A with firmness n between t and t', then the posterior doxastic state of s at t' is characterized by the $A \rightarrow n$ conditionalization of κ, $\kappa_{A\rightarrow n}$.

Since both prior and posterior doxastic state are represented by ranking functions, this law of conditionalization is iterable. As should be clear, if the agent does not believe that A, then conditionalization by A with firmness $n > 0$ corresponds to AGM revision by A. In particular, the agent's posterior *belief set* will be the same no matter which $n > 0$ is chosen. Now, if the agent *does* believe that A, then conditionalization by A with firmness $n = 0$ corresponds to AGM contraction by A. In particular, after conditionalization $\kappa_{A\rightarrow 0}(A) = \kappa_{A\rightarrow 0}(\bar{A}) = 0$. Notice that since n must be 0 for contraction to occur, the agent's posterior *doxastic state* depends only on the proposition being contracted from the agent's belief set.[11]

Now, there is much more about Ranking Theory that is worth discussing, not the least of which is how it compares to a standard Bayesian probabilistic framework with respect to such topics as confirmation, independence, decision-making, etc., but such discussion lies beyond the intended scope of this paper.[12] Instead, let's consider a potential problem for Ranking Theory. I claimed above that *differences* between ranks are meaningful, so that if some possibilities are assigned a rank of 3 and others are assigned a rank of 4, this corresponds to a doxastic state that is different from the state in which the latter possibilities are assigned a rank of 5 instead, *even if* in this second state there are no possibilities assigned a rank of 4. However, I gave no indication of *what* such differences correspond to, nor how to *measure* such differences. To put the point another way, I

[11]If the agent believes that A, then $\kappa(\bar{A}) = m$ for some $m > 0$, and $A \rightarrow n$ conditionalization with $n > 0$ will result in no change to the agent's belief set, but the ranks of the $\bar{A}$-possibilities will *increase* if $n > m$ (the information increased the agent's disbelief in certain propositions), *decrease* if $n < m$ (the information decreased the agent's disbelief in certain propositions), and remain the same if $n = m$. If the agent does not believe that A, then $\kappa(\bar{A}) = 0$, and $A \rightarrow 0$ conditionalization will result in $\kappa_{A\rightarrow 0}(A) = 0$, so if $\kappa(A) > 0$ then $A \rightarrow 0$ conditionalization will in fact result in contraction of the agent's belief that $\bar{A}$. So, $A \rightarrow 0$ conditionalization is only AGM contraction by A if $\kappa(A) = 0$, and in order to properly define the contraction $\kappa_{\div A}$ of κ let's stipulate that $\kappa_{\div A} = \kappa$ if $\kappa(\bar{A}) = 0$.

[12]Once again, I refer the interested reader to (Spohn, unpublished).

gave no *operational definition* of ranks. It is one thing to claim that the agent considers possibility w_1 more implausible than possibility w_2; it is quite another to claim that the agent considers w_1 to be implausible with rank 14 and w_2 with rank 7, so w_1 is ranked exactly 7 higher than w_2. Furthermore, given nothing but an agent's belief set, one can determine the agent's entire ordering of disbelief by comparing the agent's different posterior belief sets under different hypothetical contractions, provided such contractions satisfy the AGM constraints we discussed above. More carefully, in the AGM framework B is at least as disbelieved as A exactly if either B is maximally disbelieved or after contracting by $\bar{A} \wedge \bar{B}$ the agent does not believe that $\bar{A}$ (Gärdenfors, 1988). So, let's define $\trianglelefteq_{\div}$, the ordering generated by $\div$, such that for all $A, B \in \mathcal{A}$, $A \trianglelefteq_{\div} B$ iff either $B = \emptyset$ or $\div(\bar{A} \cap \bar{B}) \nsubseteq \bar{A}$.

What about the Ranking Theorist? Can the Ranking Theorist give an operational definition of ranks in terms of nothing more than belief sets and hypothetical contractions? According to Matthias Hild and Wolfgang Spohn (2008), the answer is yes. Hild and Spohn have demonstrated that given an agent's belief set, one can determine the agent's entire *ranking function* (up to multiplication by a constant) by comparing the agent's different posterior belief sets under different hypothetical *iterated* contractions (at most three successive hypothetical contractions are needed in each case), provided such iterated contractions satisfy six axioms given by Hild and Spohn.[13] Hild and Spohn demonstrate this by showing how one can use an iterated contraction function characterizing an agent's doxastic state (and satisfying their six axioms) to construct a difference comparison of

[13]In (Hild and Spohn, 2008), Hild and Spohn weaken the definition of a ranking function so that the range of the function doesn't have to be well-ordered. In particular, Hild and Spohn define ranking functions as maps from propositions to nonnegative real numbers (rather than from possibilities to natural numbers) and add the law of finite disjunction $\kappa(A \cup B) = \min\{\kappa(A), \kappa(B)\}$ as part of the definition. Hild and Spohn claim this weaker definition facilitates the connection with measurement theory, since measurement scales usually consist of real numbers. I have chosen to stick with the stronger definition of a ranking function in order to elucidate the *motivation* for Ranking Theory in terms of Belief Revision, as well as to capture the structure of Ranking Theory in an intuitive way without unnecessary complication. Additionally, if one uses the weaker definition, several other definitions will also have to be changed accordingly. See (Hild and Spohn, 2008) for details.

pairs of propositions in terms of relative implausibility (i.e., the difference in implausibility between A and B is greater than the difference in implausibility between C and D). They then use a theorem of Krantz, Luce, Suppes, and Tversky (1971) in order to prove the existence of a ranking function that satisfies these difference comparisons and is unique up to multiplication by a constant. Hence, iterated contraction functions satisfying their six axioms provide a measurement of ranks on what is known as a *ratio* scale.

Now, Hild and Spohn admit that such difference comparisons are not intuitively well accessible. However, (at most) three successive hypothetical contractions of a belief set are much more accessible, and since one can construct the needed difference comparison from such hypothetical iterated contractions (provided they satisfy the axioms), this seems to answer the operationalist critic. Ranks correspond to dispositions to successively contract one's current belief set, when necessary, in accordance with certain axioms. As should be clear, this measurement result is only as valuable as the independent plausibility of the axioms it is based upon. In particular, we gave independent motivation above for the constraints on an AGM contraction function, and hence the fact that one can construct an ordering of disbelief uniquely via a belief set and hypothetical single AGM contractions can be thought of as operationalist justification for our use of orderings of disbelief.[14] Similarly, the six axioms for *iterated* contraction ought to have strong intuitive justification. Let's briefly consider the six axioms (modified to reflect our previous definitions) and the justification provided by Hild and Spohn.

Let $\div$ be a function from the set of all finite sequences of propositions from $\mathcal{A} - \{W\}$ to the set of all belief sets made up of propositions from $\mathcal{A}$ and let $\div_{\langle S \rangle}$ be the function assigning the value $\div\langle S, S' \rangle$ to each finite sequence S' of propositions from $\mathcal{A} - \{W\}$.

[14] Alternatively, if one finds orderings of disbelief intuitively accessible, one could think of such orderings of disbelief as justifying the AGM axioms for contraction, since one can construct a unique AGM contraction function from such an ordering. This move doesn't seem available to the Ranking Theorist, however, since as I claimed above, it is difficult to see what ranks of disbelief are supposed to be at an intuitive level besides simply positions in an ordering of disbelief. So, one would need strong *independent* justification of the full Ranking-Theoretic doxastic structure in order to justify the proposed axioms of iterated contraction in this way.

Then $\div$ is an *iterated contraction function* exactly if for any $A, B, C \in \mathcal{A} - \{W\}$ and any finite sequence S of propositions from $\mathcal{A} - \{W\}$:

(IC1) the function $A \mapsto \cap(\div\langle A\rangle)$ (where $\cap(\div\langle A\rangle)$ is the core of the belief set $\div\langle A\rangle$) is a propositional AGM contraction function,

(IC2) if $A \notin \div\langle W\rangle$, then $\div\langle A, S\rangle = \div\langle S\rangle$,

(IC3) if $\bar{A} \cap \bar{B} = \emptyset$, then $\div\langle A, B, S\rangle = \div\langle B, A, S\rangle$,

(IC4) if $A \subseteq B$ and $A \cup \bar{B} \notin \div\langle A\rangle$, then $\div\langle A \cup \bar{B}, B, S\rangle = \div\langle A, B, S\rangle$,

(IC5) if both $A \subseteq \bar{C}$ or $A, B \subseteq C$ and $A \trianglelefteq_{\div} B$, then $A \trianglelefteq_{\div\langle C\rangle} B$, and if the inequality in the antecedent is strict, that of the consequent is strict, too,

(IC6) $\div_{\langle S\rangle}$ is an IC.

The justification for (IC1) should be clear: iterated contractions should reduce to AGM contractions in the single case. (IC2) merely says that if A isn't even believed, then contraction by A followed by contraction by a sequence of propositions should result in the same belief set as contraction by the sequence of propositions alone. (IC6) says that the function will still satisfy the axioms at each stage of iteration (it allows Hild and Spohn to state the other axioms more clearly). (IC5) is equivalent to the widely-accepted postulates of Darwiche and Pearl (1997), which roughly correspond to the idea that contraction or revision by C should not change the ordering *within* C or *within* $\bar{C}$, but rather should only change their orderings relative to each other (see (Hild and Spohn, 2008) for details).

(IC3) corresponds to restricted commutativity. The idea is, if $\bar{A}$ and $\bar{B}$ are logically incompatible, then $\bar{B} \subseteq A$ and $\bar{A} \subseteq B$. So, if you give up your belief that A, for instance, then your belief that B should be just as entrenched as before, since contracting your belief that A corresponds to no longer disbelieving in some $\bar{A}$-possibilities while leaving one's disbelief in any A-possibility unchanged, and there are *no* $\bar{A}$-possibilities that are also $\bar{B}$-possibilities, and thus all $\bar{B}$-possibilities should remain unchanged. Therefore, there shouldn't be any interaction between giving up disbelief in $\bar{A}$ and giving up disbelief in $\bar{B}$, and hence the order of contraction shouldn't matter.

(IC4) is known as path independence. Hild and Spohn motivate it as follows: Suppose you believe that A and B is a logical consequence of A, so you also believe that B and believe that $A \vee \bar{B}$. Now, if you were to give up your belief that A, you would also need to give up either your belief that B or your belief that $A \vee \bar{B}$. Suppose that you would give up your belief that $A \vee \bar{B}$ in such a case. Now, (IC4) says that given these assumptions, it wouldn't matter whether you did in fact contract A and then B or you just went ahead and contracted $A \vee \bar{B}$ directly and then B; you would end up in the same doxastic state in either case.

So, the six axioms seem well-justified. Before considering Stalnaker's counterexample, let's address a final worry. Hild and Spohn (2008) demonstrate that one can take an iterated contraction function that satisfies (IC1)-(IC6) and use it to construct a ranking function that is unique up to multiplication by a constant. They also demonstrate that, given a ranking function, one can construct a unique iterated contraction function (since $A \to 0$ conditionalization is iterable and the contraction of a ranking function κ by a sequence of propositions corresponds to propositional AGM contraction at each step, so long as we stipulate that $\kappa_{\div A} = \kappa$ whenever $\kappa(\bar{A}) = 0$). However, they never demonstrate that the iterated contraction function induced by a ranking function will *itself* statisfy (IC1)-(IC6)! If this were *not* the case, then any intuitive justification for (IC1)-(IC6) would be justification for *rejecting* the Ranking-Theoretic framework as it currently stands, and hence Hild and Spohn's argument would be self-defeating. Thankfully, this lacuna can be filled, and in (Bice, 2008) I proved that the iterated contraction function induced by a ranking function will always satisfy (IC1)-(IC6).

4 Stalnaker's Counterexample

Let's consider Robert Stalnaker's (2009) counterexample to Ranking-Theoretic iterated belief revision. I will quote the counterexample in full, and then proceed to explain why the counterexample fails. Here it is:

"Consider this example: Fair coins are flipped in each of two rooms. Alice and Bert (who I initially take to be reliable) report to me, independently, about the results: Alice tells me that the coin in room A came up heads, while Bert tells me that the coin in room B came up heads, and so this is what I believe at stage one. Because my sources were independent, my belief revision policies, at stage one, will give priority to the $H_A T_B$ and $T_A H_B$ possibilities over the $T_A T_B$ possibility (Were I to learn that Bert was wrong, I would continue to believe that Alice was right, and vice versa). But now Carla and Dora, also two independent witnesses whose reliability, in my view, trumps that of Alice and Bert, give me information that conflicts with what I heard from Alice and Bert. Carla tells me that the coin in room A came up tails, and Dora tells me the same about the coin in room B. These two reports are also given independently, though we may assume simultaneously. This is stage two. Finally (stage three), Elmer, whose reliability trumps everyone else, tells me that that the coin in room A in fact landed heads (So Alice was right after all). What should I now believe about the coin in room B? DP's postulate (C2) requires that I return to the original belief that the coin in room B came up heads. Even though Dora's information had overturned Bert's information about the second coin, and even though Elmer provided no information at all about the result of the coin flip in room B (we may assume he knew nothing about it), Elmer's information still forces us to change our belief about this result (if we follow the DP constraint, C2)."

Stalnaker considers this to be a counterexample to the Darwiche-Pearl revision postulate (C2), which is roughly: if $A \subseteq \bar{B}$, then revision by B followed by revision by A should result in the same belief set as revision by A alone. This is equivalent to requiring that the ordering of disbelief within $\bar{B}$ not change after revision by B. Stalnaker points out that the Ranking Theorist is committed to the Darwiche-Pearl postulates, and indeed one can easily show that every ranking function will satisfy (C2). What should the Ranking Theorist say in response to this example?

The first thing to point out is that, as stated, this example cannot be formulated within a Ranking-Theoretic framework. To see this, let's begin by considering a simple algebra $\mathcal{A}$ over a set of possibilities W. Let W consist of four possibilities corresponding to the propositions $A \cap B$, $A \cap \bar{B}$, $\bar{A} \cap B$, and $\bar{A} \cap \bar{B}$, where A is the proposition that the coin in room A landed heads, while $\bar{A}$ is the proposition that the coin in room A landed tails (similarly for B and $\bar{B}$). As Stalnaker has presented the example, the agent begins with no information about how each of the coins (each of which the agent knows to be fair) have landed, and hence let's represent the agent as initially suspending judgment. So the ranking function κ that characterizes the agent's doxastic state at stage zero is such that $\kappa(A \cap B) = \kappa(A \cap \bar{B}) = \kappa(\bar{A} \cap B) = \kappa(\bar{A} \cap \bar{B}) = 0$.

Now, Stalnaker represents the initial revision at stage one as a revision by $A \cap B$.[15] However, since Alice and Bert are independent and reliable witnesses, Stalnaker wants it to be the case that, after revision, the agent considers the $A \cap \bar{B}$ and $\bar{A} \cap B$ possibilities to be more plausible than the $\bar{A} \cap \bar{B}$ possibility, since if the agent were to learn that one of the two witnesses was wrong, she would continue to believe that the other witness was right. But consider how revision by $A \cap B$ is handled in a Ranking-Theoretic framework: since Alice and Bert are considered reliable witnesses, the agent conditionalizes on $A \cap B$ with firmness $n > 0$ for some natural number n. So, after $A \cap B \rightarrow n$ conditionalization, $\kappa_{A \cap B \rightarrow n}(A \cap B) = \kappa(A \cap B | A \cap B) = \kappa(A \cap B) - \kappa(A \cap B) = 0$, while $\kappa_{A \cap B \rightarrow n}(\bar{A} \cup \bar{B}) = \kappa(\bar{A} \cup \bar{B} | \bar{A} \cup \bar{B}) + n = n$, and so the agent believes that both coins landed heads, as desired. However, $\kappa_{A \cap B \rightarrow n}(A \cap \bar{B}) = \kappa_{A \cap B \rightarrow n}(\bar{A} \cap B) = \kappa_{A \cap B \rightarrow n}(\bar{A} \cap \bar{B}) = n$, since all are subsets of $\bar{A} \cup \bar{B}$ and hence their ranks will all rise from 0 to n. Hence, the agent treats the other three possibilities as equally implausible, contrary to what Stalnaker wants.[16]

What is going wrong here? In our Ranking-Theoretic framework, conditionalizing on

[15] Stalnaker (2009, footnote 31).

[16] Should we abandon our assumption that the agent begins by suspending judgment? In order to get the desired asymmetry between the two-tails possibility $\bar{A} \cap \bar{B}$ and the one-heads possibilities $A \cap \bar{B}$ and $\bar{A} \cap B$ after the first revision, we would need to assume that the agent starts out disbelieving that both coins landed tails to a greater degree than her degree of disbelief in either of the possibilities in which one coin landed heads and one coin landed tails. But this is clearly not what Stalnaker intended. Furthermore, if we were to follow this example to its conclusion, although the agent would end stage 3 by switching her belief about the coin in room B from $\bar{B}$ to B, this conclusion would not be counterintuitive for basically the same reasons I give regarding case 2 below.

the proposition $A \cap B$ with some firmness n corresponds to accepting (with firmness n) the information that the state of the world is such that both coins have landed heads. Since the agent initally suspended judgment, after accepting this information, the agent will consider the other three possible states of the world to be equally implausible. Stalnaker wants to model a doxastic state in which the agent receives two *independent* pieces of information about the state of the world; namely, that the coin in room A landed heads, and that the coin in room B landed heads. In order to model this doxastic state in our Ranking-Theoretic framework, we should treat the agent's acceptance of these two pieces of information as *separate* revisions. So, let's consider two modified forms of the example: in the first, standard case, acceptance of each independent piece of information will be modeled by a separate revision; in the second case, we will consider what happens when the agent simply revises by the conjunctions $A \cap B$ and then $\bar{A} \cap \bar{B}$, as Stalnaker originally intended.[17]

Case 1: Let $\kappa' = \kappa_{A \to n}$. So, after Alice informs the agent that the coin in room A landed heads, which the agent accepts with firmness n, the agent's posterior doxastic state is characterized by the ranking function κ'. In particular, $\kappa'(A \cap B) = \kappa'(A \cap \bar{B}) = 0$ and $\kappa'(\bar{A} \cap B) = \kappa'(\bar{A} \cap \bar{B}) = n$. So, the agent believes that coin A landed heads but suspends judgment regarding coin B. Now, let $\kappa'' = \kappa'_{B \to n}$. Now, $\kappa''(A \cap B) = 0$ while $\kappa''(A \cap \bar{B}) = \kappa''(\bar{A} \cap B) = n$ and $\kappa''(\bar{A} \cap \bar{B}) = 2n$. So, after the second revision, the agent believes that both coins A and B landed heads and considers the possibilities in which one coin landed heads and one coin landed tails more plausible than the possibility in which

[17] Could we simply enrich W so that we can express the propositions that Alice said that the coin in room A landed heads and that Bert said that the coin in room B landed heads as subsets of W, and then consider their conjunction? Perhaps, but we would no longer have a counterexample to (C2), since the conclusion that Stalnaker wants would no longer follow from (C2). Furthermore, one could then model Stalnaker's intuitions regarding this example by including information about the perceived reliability of particular witnesses in the agent's prior doxastic state in a straightforward way (for example: $\kappa(\bar{A}|\text{Alice said } A)$). Similar remarks apply to other attempts to enrich the possibility space. Stalnaker himself expressed the information accepted as simply whether coin A (or B) landed heads or tails (2009, footnote 31), as is needed to invoke (C2).

both coins landed tails. Notice in particular that the agent would have been in precisely the same doxastic state if she had first revised by B and then A. Let $\kappa''' = \kappa''_{\bar{A}\to n+m}$, for some natural number m (with $n+m$ roughly corresponding to the perceived reliability of Carla). So, $\kappa'''(A \cap B) = n+m$, $\kappa'''(A \cap \bar{B}) = 2n+m$, $\kappa'''(\bar{A} \cap B) = n-n = 0$, and $\kappa'''(\bar{A} \cap \bar{B}) = 2n-n = n$. Similarly, let $\kappa^{(4)} = \kappa'''_{\bar{B}\to n+m}$. So, $\kappa^{(4)}(A \cap B) = 2n+2m$, $\kappa^{(4)}(A \cap \bar{B}) = 2n+m-n = n+m$, $\kappa^{(4)}(\bar{A} \cap B) = n+m$, and $\kappa^{(4)}(\bar{A} \cap \bar{B}) = n-n = 0$. So, at the end of stage two, the agent believes that both coins landed tails and considers the possibility in which both coins landed heads much more implausible than either of the two possibilities in which one coin landed heads and one coin landed tails. Again, the order in which the agent revises by $\bar{A}$ and $\bar{B}$ is irrelevant. In the final stage, Elmer informs the agent that the coin in room A landed heads. Let $\kappa^{(5)} = \kappa^{(4)}_{A\to n+m+s}$, for some natural number s (with $n+m+s$ roughly corresponding to the perceived reliabilitiy of Elmer). So, $\kappa^{(5)}(A \cap B) = 2n+2m-(n+m) = n+m$, $\kappa^{(5)}(A \cap \bar{B}) = n+m-(n+m) = 0$, $\kappa^{(5)}(\bar{A} \cap B) = n+m+(n+m+s) = 2n+2m+s$, and $\kappa^{(5)}(\bar{A} \cap \bar{B}) = 0+(n+m+s) = n+m+s$.

Thus, after the agent accepts Elmer's information with firmness $n+m+s$, she believes that both the coin in room A landed heads and the coin in room B landed tails, which is precisely the intuitive conclusion that Stalnaker claimed is correct. So, the agent does *not* switch her belief regarding coin B, and hence the counterexample fails in this case.

Case 2: Let $\kappa' = \kappa_{A\cap B\to n}$. As stated above, after stage one, $\kappa'(A \cap B) = 0$ and $\kappa'(A \cap \bar{B}) = \kappa'(\bar{A} \cap B) = \kappa'(\bar{A} \cap \bar{B}) = n$. So, after stage one, the agent believes that the state of the world is such that both coins landed heads, and considers the other three possibilities equally implausible. Let $\kappa'' = \kappa'_{\bar{A}\cap\bar{B}\to n+m}$. So, $\kappa''(A \cap B) = n+m$, $\kappa''(A \cap \bar{B}) = \kappa''(\bar{A} \cap B) = 2n+m$, and $\kappa''(\bar{A} \cap \bar{B}) = 0$. So, at stage two, the agent receives and accepts the information that the world is in the tails/tails state, and hence after revising her beliefs, she believes that the tails/tails state obtains but considers the heads/heads state more plausible than either the heads/tails state or the tails/heads state, since the heads/heads state was previously the most plausible state. Finally, let $\kappa''' = \kappa''_{A\to n+m+s}$. So, $\kappa'''(A \cap B) = n+m-(n+m) = 0$, $\kappa'''(A \cap \bar{B}) = 2n+m-(n+m) = n$, $\kappa'''(\bar{A} \cap B) = 2n+m+(n+m+s) = 3n+2m+s$, and $\kappa'''(\bar{A} \cap \bar{B}) = n+m+s$.

Thus, after the agent accepts the information that the state of the world is such that the coin in room A landed heads, she revises her belief that the coin in room B landed tails to the belief that the coin in room B also landed heads, *precisely because* she considered the heads/tails state much more implausible than the heads/heads state. In case 2, the conclusion that Stalnaker wants does in fact obtain, but it is no longer counterintuitive, since we are no longer modeling the example that Stalnaker intended. We have modeled a case in which the agent starts out suspending judgment, receives good information that the state of the world is heads/heads, receives better information that the state of the world is tails/tails, and finally receives even better information that the state of the world is either heads/heads or heads/tails. In this case, there is no *interaction* between states of the world in the sense that evidence that one state of the world obtains makes other states of the world more plausible to the agent; all four states of the world are mutually exclusive and exhaustive. In particular, after revising by the proposition that the state of the world is heads/heads, the agent does not consider the heads/tails or tails/heads possibilities more plausible than the tails/tails possibility, since she hasn't received any information regarding their relationship, but rather merely received the information that none of them obtain. This is clearly not the appropriate way to model Stalnaker's example.

Thus, Stalnaker's counterexample to Ranking-Theoretic iterated belief revision fails. Therefore, given our discussion above regarding the limitations of AGM belief revision and the strengths of Ranking-Theoretic belief revision, it seems we have strong reasons for accepting the additional structure provided by the Ranking-Theoretic approach. Furthermore, given the measurement result of Hild and Spohn (2008) and the intuitive plausibility of their six axioms, this additional structure can be given operational justification. Thus, Ranking Theory seems to be the clear choice for modeling rational constraints on revising non-quantitative states of belief.

References

[1] Alchourron, Carlos E., Peter Gärdenfors, David Makinson (1985), "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", *Journal of Symbolic Logic* 50, 510-530.

[2] Bice, Nathan (2008), "Old Evidence, Belief Revision, and Ranking Theory", Senior Honors Thesis, University of Wisconsin.

[3] Darwiche, Adnan, Judea Pearl (1997), "On the Logic of Iterated Belief Revision", *Artificial Intelligence* 89, 1-29.

[4] Gärdenfors, Peter (1988), *Knowledge in Flux. Modeling the Dynamics of Epistemic States.* Cambridge, MA: MIT Press.

[5] Harper, William (1977), "Rational Conceptual Change", in: *PSA 1976*, East Lansing, MI: Philosophy of Science Association, vol. 2, pp. 462-494.

[6] Hild, Matthias, Wolfgang Spohn (2008), "The Measurement of Ranks and the Laws of Iterated Contraction". *Artificial Intelligence*, 172: 1195-1218.

[7] Krantz, David, R. Duncan Luce, Patrick Suppes, Amos Tversky (1971), *Foundations of Measurement, Vol. I.* New York: Academic Press.

[8] Levi, Isaac (1977), "Subjunctives, Dispositions and Chances", *Synthese* 34: 423-455.

[9] Makinson, David (1987), "On the Status of the Postulate of Recovery in the Logic of Theory Change", *Journal of Philosophical Logic* 16: 383-394.

[10] Spohn, Wolfgang (1988), "Ordinal Conditional Functions. A Dynamic Theory of Epistemic States", in: W. L. Harper, B. Skyrms (eds.), *Causation in Decision, Belief Change, and Statistics, vol. II*, Dordrecht: Kluwer, pp. 105-134.

[11] Spohn, Wolfgang (unpublished), *Ranking Theory: A Tool for Epistemology.* manuscript.

[12] Stalnaker, Robert (2009), “Iterated Belief Revision”, *Erkenntnis* 70: 189-209.