Notre Dame Journal of Formal Logic
Volume 51, Number 1, 2010

On the Ramsey Test without Triviality

Hannes Leitgeb

Abstract We present a way of classifying the logically possible ways out of Gärdenfors' inconsistency or triviality result on belief revision with conditionals. For one of these ways—conditionals which are not descriptive but which only have an inferential role as being given by the Ramsey test—we determine which of the assumptions in three different versions of Gärdenfors' theorem turn out to be false. This is done by constructing ranked models in which such Ramsey-test conditionals are evaluated and which are subject to natural postulates on belief revision and acceptability sets for conditionals. Along the way we show that in contrast with what Gärdenfors himself proposed, there is no dichotomy of the form: either the Ramsey test has to be given up or the Preservation condition. Instead, both of them follow from our postulates.

1 Introduction

This paper[1] deals with conditionals $A > B$—and only with such—whose antecedent A and consequent B speak about the "outside world" (rather than, e.g., some agent's state of mind) and where neither A nor B itself includes the nonmaterial if-then symbol '$>$'; call such conditionals 'simple conditionals'.[2] Conditionals of this kind have potentially two roles to play in our reasoning and communicative behavior:

(i) they may track our inferential dispositions to (hypothetically) change some of our beliefs about the world given certain suppositions about the world;
(ii) they may describe the world to be so and so; for instance, they may track the world's disposition to change some of its worldly states of affairs given some other of its worldly states of affairs are changed.

Both of these roles (i) and (ii) can be made precise either on a qualitative-ordinal scale or on a quantitative-probabilistic one; here we will concentrate just on the former (although there are probabilistic counterparts to all of the positions and moves that we are going to discuss).

Received December 7, 2009; printed April 26, 2010
2010 Mathematics Subject Classification: Primary, 03A05
Keywords: Ramsey test, conditionals, belief revision, acceptability
 10.1215/00294527-2010-003

The so-called Ramsey test for conditionals, which was suggested by Frank Plumpton Ramsey in 1929 (see [18]), conveys an evaluation procedure for conditional sentences that is capturing their role (i). Later on, Stalnaker [24, p. 102], generalized Ramsey's original formulation to include also cases in which the antecedent is believed to be false; in his words,

> This is how to evaluate a conditional: First, add the antecedent (hypothetically) to your stock of beliefs; second, make whatever adjustments are required to maintain consistency (without modifying the hypothetical belief in the antecedent); finally, consider whether or not the consequent is then true.

This test for the acceptance of conditionals was formalized later still by Gärdenfors [4; 5] in the language of qualitative belief revision and then questioned on grounds of inconsistency. This is Gärdenfors' version of the Ramsey test:

$$B \in G * A \text{ iff } A > B \in G.$$

Here G is an arbitrary belief set of an agent that is meant to include all sentences believed true by that agent at time t, and $*$ is the same agent's so-called belief revision operator which represents the relevant suppositional aspects of the agent's inferential dispositions at t. So '$B \in G * A$' says that supposing A at t leads the agent to revise her beliefs in the way that she ends up with $G * A$ as her new (hypothetical) belief set under that supposition, and '$A > B \in G$' expresses that $A > B$ is believed true according to G.[3]

If this variant of the Ramsey test is added to some of the standard rationality principles for belief revision operators $*$ (the AGM postulates, cf. Alchourrón, Gärdenfors, and Makinson [2], Gärdenfors [5]), and if $*$ is assumed to be nontrivial, a logical contradiction can be derived; only trivialized revision operators (or revision operators defined on a trivialized set of belief sets) allow for the Ramsey test in the form suggested by Gärdenfors. In other words,

$$\frac{\begin{array}{l}\text{AGM Axioms of Belief Revision}\\ \text{Gärdenfors' Ramsey test}\\ \text{Nontriviality Assumptions}\end{array}}{\bot}$$

or equivalently

$$\frac{\begin{array}{l}\text{AGM Axioms of Belief Revision}\\ \text{Gärdenfors' Ramsey test}\end{array}}{\text{Triviality}}$$

can be shown to be logically valid; hence we will keep referring to this as an inconsistency or triviality result. The interpretations of it are legion: Gärdenfors' own conclusion from his inconsistency or triviality result is formulated cautiously: as he expresses it in Gärdenfors [4, p. 86], his "present position" is that the Ramsey is the culprit of the problem. Others have disagreed.[4] Overall, the reception of the result has been patchy and even slightly chaotic, for the simple reason that the distinction between (i) and (ii) above is quite often not being taken into account.

The aim of this article is to clarify the situation by explicit reference to the two potential roles of conditionals as explained above and to uncover how so-called Ramsey-test conditionals[5] which play role (i) but not role (ii) manage to avoid inconsistency or triviality in the face of Gärdenfors' theorem. As far as the former task is concerned, to which we will turn in Section 2, we do not claim that we will able to cover the complete territory of logically possible positions vis-à-vis Gärdenfors' result (in whatever sense of completeness). For example, we will neglect the possibility that a conditional might express different propositions in different pragmatic or epistemic contexts (as suggested by Lindström [15]). But we hope Section 2 will still be general enough in order to see what the basic lines of reasoning are that may lead us out of the unfortunate conclusions of the arguments above. Afterward we will turn to our second task of following one such line of reasoning in detail, that is, determining how Ramsey-test conditionals manage to avoid inconsistency and triviality. This is the corresponding plan of the paper from Section 3: in Section 3 we will outline the formal background framework of the models in which we will evaluate Ramsey-test conditionals. Section 4 is devoted to state some minimal requirements on belief revision operators; they will all be postulates which have counterparts among the axioms of AGM (or conclusions from them), but the requirements will now involve Ramsey-test conditionals and so-called acceptability sets unlike the original AGM postulates. We will show that the Ramsey test in a version in which acceptability replaces belief is actually following from these minimal requirements rather than contradicting them. In Section 5, we will suggest one possible manner of extending the minimal requirements on belief revision according to an underlying maxim of minimal mutilation, which will give us a version of the so-called Preservation principle that will apply even to conditionals. While these stronger principles are still consistent with the version of the Ramsey test that we will be dealing with, and while they conform to what many critics of Stalnaker's general version of the Ramsey test have taken for granted themselves, they contradict some of the auxiliary hypotheses that have been employed in order to prove the failure of Gärdenfors' variant of the Ramsey test; it is these assumptions which have to be given up in the case of Ramsey-test conditionals as they do not carry over from belief revision for statements which express propositions to conditionals that do not express propositions. Along the way, it will also follow that there is no dichotomy of the form "either the Ramsey test has to go or Preservation must" as Gärdenfors [5] seemed to suggest. We will demonstrate these results by turning to three representative presentations of Gärdenfors' original inconsistency or triviality result: Lindström and Rabinowicz [16] presuppose what they call the "Nontriviality" condition; however, as we will show, the falsity of this condition in a certain epistemic situation does not mean that the situation is implausible or trivial. In fact, the Nontriviality condition cannot be satisfied for Ramsey-test conditionals and the acceptability sets in which they are included, as we will show in the corresponding Section 6. Section 7 will be devoted to Gärdenfors' own exposition of his result in [5], for which he assumes additional revision principles which are justified as long as one is only interested in propositional outputs of belief revision but which cease to be justified if the result of a belief revision is supposed to determine an acceptability set which includes as well conditionals that fail to express propositions. Segerberg's [22] version of Gärdenfors' theorem, which we will discuss in Section 8, suffers from a similar "problem" once conditionals have been assumed to be Ramsey-test conditionals, as it is the case in our toy model. Each of the three versions of the triviality result that we will consider constitutes a logically valid, interesting, and valuable argument in itself; however, our analysis is going to undermine the relevance of these arguments for Ramsey-test conditionals. Section 9 will summarize what has been achieved. The concluding appendix complements our discussion by laying down some of the relevant logical background notions.

2 Cleaning House

Let us start by giving an unduly trivial answer to the question, "Which of Gärdenfors' premises in the arguments of Section 1 must be given up in order to circumvent the problematic conclusions of these arguments?" Since the AGM axioms of belief revision are *constitutive* of belief revision in the AGM-sense of the word, of course Gärdenfors' variant of the Ramsey test is the one to be dropped (as nontriviality cannot be the bad guy). End of story.

But that question is, of course, completely beside the point: in order to see why and what the real question is, we start by reformulating the arguments from Section 1 as follows. Let Op be any epistemic operation that takes two arguments—a belief set G and a formula A—and which maps them to another belief set $Op(G, A)$. Reformulate the AGM axioms of belief revision in the way that every occurrence of the form '$G * A$' is replaced by '$Op(G, A)$'; the resulting axioms express constraints on Op which may or may not be true, but we no longer consider them to be constitutive of the operation Op. We apply the same syntactic procedure to the original nontriviality assumptions on $*$ which thus become additional assumptions on Op. Finally, we reformulate Stalnaker's informal description of the Ramsey test as quoted in Section 1 in slightly more formal terms but also in a way that is less committing than Gärdenfors' original formulation: for the sake of simplicity, we restrict ourselves to arbitrary simple conditionals $A > B$ again; let G be an agent's belief set at time t; then we take the Ramsey test for Op to constrain the epistemic operator Op such that it conforms to

$$B \in Op(G, A) \text{ iff } A > B \text{ is acceptable relative to } G \text{ (and } Op).$$

What this expresses is that $Op(G, A)$ takes over the role of the hypothetical belief set that results from modifying the agent's actual belief set G at t by supposing A, for which it is then claimed that it includes B as a member if and only if $A > B$ is acceptable to the agent relative to the belief set G (and the given operator Op). If Op satisfies the Ramsey test in this version, it may therefore be expected to be some kind of supposition operator. We leave open for the moment whether 'is acceptable relative to the belief set G' means the same as 'is a member of the belief set G' or not, or whether they at least have the same extension; this will be the crucial point in Section 2.2 on Ramsey-test conditionals. Note that by our restriction to simple conditionals $A > B$, we will never instantiate the Ramsey test such that A above is itself a conditional—supposing a conditional, that is—nor such that B above is a conditional—that is, where a conditional would be demanded to be among the outcomes of a supposition by the Ramsey test.

Given these reformulations, we are finally in the position to restate the arguments from Section 1 in the following form:

AGM Axioms of Belief Revision applied to Op
Ramsey test for Op
Nontriviality Assumptions on Op
____
$\bot$

and

AGM Axioms of Belief Revision applied to Op
Ramsey test for Op
____
Triviality for Op

If the right-hand side of the Ramsey test for *Op* is understood to be extensionally equivalent to '$A > B$ is a member of the belief set G', that is, to '$A > B \in G$', then by Gärdenfors' theorem these two arguments are logically valid again. On the other hand, it is well known—and it is also going to follow from the model construction from Section 3 of this paper—that if this very step is not being taken, then the arguments turn out not to be valid anymore. In any case, we are now free to question each of the premises in the arguments dependent on any choice of any epistemic operation *Op* that we please, depending on how we intend to interpret 'is acceptable relative to the belief set G (and *Op*)', and, of course, relative to the intended interpretation of the conditional sign '$>$'. Depending on these choices, it might be the case that the AGM axioms are no longer satisfied by *Op*, the resulting version of the Ramsey test might be implausible, and even the so-called nontriviality requirements on *Op*—whose name got preserved from the original formulation of the first argument for $*$ in Section 1—might actually cease to express nontriviality when an epistemic operator other than any AGM belief revision operator is employed. So the real question at this point is, *Which possibilities are there to choose these parameters such that the problematic conclusions of these arguments do not follow logically from their premises?* In order to answer this question on a systematic basis we need to say more about the two roles that conditionals can play as sketched in Section 1. There are three possible cases to consider.

2.1 The first case: role (ii), not role (i) We begin by considering conditionals $A > B$ which are *playing role (ii) but not role (i)*: By playing role (ii), any such conditional $A > B$ must be true or false at a possible world, depending on whether the world is as described by the conditional; for instance, depending on whether that world has, in fact, the disposition to change its worldly states such that B is the case given the world is changed such that A is the case, or whatever other semantic (rather than pragmatic) meaning is conveyed by $>$. As usual, the set of possible worlds in which $A > B$ is true may be regarded as the proposition that is expressed by that conditional. Believing such a conditional to be true is then not different *in kind* from believing A to be true or believing B to be true. Indeed, if we take beliefs to be sensitive only to propositions rather than to the sentences that express them, then if such a conditional $A > B$ is believed by an agent, the content of this belief does not even involve any conditional structure at all; the belief in question is directed toward a particular set of possible worlds just as any other belief is. Any belief set G may thus include such a conditional $A > B$ in the very same sense as it might perhaps include A or B. In a similar way, also the axioms of belief revision operators $*$ in the sense of AGM are sensitive only to the level of propositions: if A and A' express the same proposition, and if the set of worlds in which all the members of G are true is identical to the set of worlds in which all the members of G' are true, then it follows from AGM's standard rationality postulates on $*$ that $G * A = G' * A'$. Accordingly, Grove's [7] representation theorem reconstructs any AGM-type belief revision operator $*$ in terms of a spheres system or linear preorder of possible worlds, and the least sphere or rank in any such semantic structure corresponds to the set of worlds in which all and only the members of the agent's current belief set are true. Indeed, even though this may not have been emphasized enough or even realized by AGM, their rationality postulates were really designed to apply to (belief sets of) formulas which are true or false and which express propositions. Since in the case that we are considering at present, conditionals $A > B$ do express propositions, the axioms of AGM belief revision may be applied to them and to the belief sets that include them as members in exactly the same sense in which they apply to nonconditional statements that express propositions. Of course, one would have to determine first the logic of $>$ from whatever the exact truth conditions of $A > B$ are like, in order to say how logical notions such as deductive closure, consistency, and logical equivalence, which figure in the AGM postulates, are to be understood for a language that includes $>$. However, there is no reason why this could not be done, and certainly the counterparts of these notions on the level of propositions are trivial (e.g., consistency translates into nonemptyness). Therefore, in the current case it is perfectly possible for conditionals $A > B$ to be members of belief sets, and it would even be possible for them to be supposed to be true such that the agent's current belief set would get changed in light of such a conditional supposition.

Among the premises on *Op* which lead to Gärdenfors' contradictory conclusion—that is, as mentioned before, the AGM axioms for *Op*, the Ramsey test for *Op*, and the nontriviality requirement for *Op*—the AGM axioms and the nontriviality assumptions should be perfectly unproblematic if we please to let $Op = *$ for some nontrivial AGM belief revision operator $*$, while the Ramsey test clearly is the one that ought to get sacrificed in this case. After all, there is no reason why the Ramsey test as stated above should apply to any of these conditionals, as they have been assumed not to play role (i) from above.

It is difficult to give any natural examples of conditionals that fall into this category, since all typical conditionals are normally assumed to play the suppositional role (i) in at least some sense, but this does not undermine the logical possibility of expressions of this kind. It is certainly not hard at all to imagine a binary logical sign $>$ which semantically forms a proposition from two argument propositions such that the axioms of belief revision apply in the standard manner to expressions of the form $A > B$ and to the belief sets that include them, but where there is no sensible epistemic operation *Op* such that *Op* and $>$ jointly satisfy the Ramsey test as formulated at the beginning of this section. In fact, it is easy to give *abstract* examples of such symbols $>$: just let $>$ be equal to any binary logical operator that cannot sensibly be regarded as forming some sort of if-then statement syntactically. For instance, if *Op* is chosen to be some nontrivial AGM belief revision operator $*$, if $>$ is chosen to be $\wedge$ or $\vee$, and if 'is acceptable relative to G' is understood as 'is a member of the belief set G', then the Ramsey test obviously comes out false.[6]

2.2 The second case: role (i), not role (ii) As far as conditionals $A > B$ are concerned which are *playing role (i) but not role (ii)*, by the absence of role (ii) such conditionals do not have truth values at worlds at all, and hence they do not express propositions either, since they do not even describe the world to be so and so at all. By role (i), any such conditional $A > B$ is acceptable for an agent in an epistemic state if and only if the agent's state satisfies the constraint that is then imposed on it through the Ramsey test; that is, $B \in Op(G, A)$. Apart from this, no formal requirement on the—as we may deduce from the Ramsey test—supposition operator *Op* follows. Since $A > B$ does not express a proposition at all, and also because neither its antecedent nor its consequent speak about the agent's epistemic states at all (by our initial assumption), the conditional cannot express *that the agent is in a particular epistemic state* but rather the acceptance of such a conditional expresses pragmatically *the epistemic state of being such that* $B \in Op(G, A)$ without actually referring to that state or describing it in any semantic sense. (Just as, say, the speech act of exclaiming "hooray" expresses pleasure but not the proposition that one is in a state of pleasure.) This is certainly very different from believing sentences such as A or B to be true. Consequently, such conditionals $A > B$ could not be members of belief sets anymore in the same sense as this proved to be possible in Section 2. Instead, it makes good sense to introduce *acceptability sets* over and above belief sets; while sentences such as A or B might be members of a belief set G, a conditional $A > B$ of the presently considered sort might be a member of the acceptability set that is determined by G (and Op) as explained by the Ramsey test: for given Op, if G is a belief set (which does not include any conditional of the form $A > B$), then $A > B$ is a member of the acceptability set that is determined by G (and Op) if and only if $B \in Op(G, A)$. We might even loosen up our terminology a bit to an extent such that the acceptability set that is determined by a belief set G is assumed to contain the belief set G as a subset. As long as it is clear which of the members of an acceptability set are believed to be true and which are accepted in the more pragmatic sense of the Ramsey test—which in the case of A, B, and $A > B$ may be established by mere syntactic inspection—this simplifying move will not cause any trouble. There is no a priori reason to believe of any of the AGM axioms that they would still apply to conditionals of that kind and to the acceptability sets in which they are included: as explained before, these axioms were designed—whether intentionally or unintentionally—for propositional purposes. It would be left to a case-by-case analysis to determine which of the AGM postulates could still be postulated to hold reasonably for such conditionals and acceptability sets, without running into any inconsistency or triviality results again. For instance, would it be possible to postulate that if an acceptability set is revised by a sentence A, such that A is consistent with all the members of that acceptability set that express propositions, then that acceptability set is a subset of the set that results from the revision? (We will answer this question affirmatively in Section 5 below.) Furthermore, for some of the postulates that one would like to preserve from AGM one would first have to specify the logic of $>$ again. In this case, the logic could not be derived from any truth conditions for $>$ as there are none; instead the logic of $>$ would have to be determined from the acceptability conditions for $>$. For instance, if Op is some AGM belief revision operator, then $A > B$ might be called satisfiable if and only if there exists some AGM belief revision operator $*$ (which could be Op itself or some other) such that $B \in G * A$; this semantic property could then be turned into a syntactic one, and so forth.[7]

Among the premises which lead to Gärdenfors' contradictory conclusion, the Ramsey test for Op (trivially, *by fiat*) is the one that ought to be preserved in this case. At least *some* of the relevant AGM postulates on Op—now formulated for acceptability sets rather than belief sets—and/or maybe even some of the nontriviality assumptions on Op, in case their justification depended on conditionals such as $A > B$ expressing propositions, will have to go, in order to avoid inconsistency or triviality again. (This will be our main topic from the next section.) Let us call conditionals of this type, which are playing role (i) but not role (ii), *Ramsey-test conditionals*. As we have seen, for such conditionals the Ramsey test takes the form

> $B \in Op(G, A)$ iff $A > B$ is acceptable relative to G (and Op)
>
> where 'is acceptable relative to the belief set G' differs in intension from 'is a member of the belief set G'.

Moreover, if $Op = *$ for some nontrivial AGM operator $*$ and hence the AGM axioms are satisfied by Op, then by Gärdenfors' theorem 'is acceptable relative to the belief set G' must also differ from 'is a member of the belief set G' *in extension*. So in this case we get

> $B \in G * A$ iff $A > B$ is acceptable relative to G (and Op)
>
> where 'is acceptable relative to the belief set G' differs in extension from 'is a member of the belief set G'.

Accordingly, if we use the terminology of acceptability sets again, we can rewrite that version of the Ramsey test for such Ramsey-test conditionals in the form $B \in G * A$ iff $A > B$ *is a member of the acceptability set that is determined by the belief set* G. If we finally forget about belief sets completely and focus instead just on acceptability sets, then we may assume that 'G' refers to some acceptability set from the start, the result of revising G by A is an acceptability set again, and we end up with

> $B \in G * A$ iff $A > B \in G$
> where G and $G * A$ are acceptability sets (rather than belief sets),

which looks exactly like Gärdenfors' original version of the Ramsey test. *However, its meaning has now changed completely in view of the fact that the conditionals in question no longer express propositions and the sets in question are acceptability sets rather than belief sets.* It is subtle but crucial changes like these which make parts of the literature on this topic so hard to digest. In any case, this last version of the Ramsey test is the one that we will study later on when we will deal with Ramsey-test conditionals in detail from Section 3. But even then we will never deal with cases in which A or B on the left-hand side of the Ramsey test above would be a conditional itself, due to our simplifying restriction to simple conditionals: A being a conditional would demand conditionals with conditional antecedents to show up on the right-hand side of the Ramsey test, which would be problematic since it would be quite unclear to understand nestings of that sort when $>$ is not truth-conditional (as in the current case). B being a conditional would need conditionals with conditional consequents to occur on the right-hand side of the Ramsey test, which could be handled in terms of iterated belief revision; for example, $C \in (G * A) * B$ iff $(B > C) \in G * A$ iff $A > (B > C) \in G$. (This is exactly what Hans Rott studies in his unpublished draft "The Ramsey test for conditionals and iterated theory change" to which we referred at the beginning of this paper.)

Turning to Ramsey-test conditionals in the sense of this section is more or less the path that, for example, Levi [13; 14] follows, in line with a purely suppositional qualitative understanding of conditionals. Dynamic-epistemic approaches to belief revision formalize the Ramsey test in a similar way (see, e.g., van Benthem [3]).

2.3 The third case: role (i), role (ii) Finally, there might be classes of conditionals $A > B$ which are *playing both role (i) and role (ii)*: these conditionals would have to be true or false at possible worlds just as the conditionals of Section 2.1, and they would therefore express propositions; also they would have to conform to the Ramsey test in the general form of the beginning of Section 2 just as the conditionals in Section 2.2 do. What Gärdenfors' theorem teaches us about them is thus the following: Either *Op* is different from any nontrivial belief revision operator $*$ in the AGM sense and, hence, either the AGM postulates do not hold for *Op* or Nontriviality is not the case (call this the *first route* out), or *Op* is a nontrivial AGM belief revision operator $*$ and 'is acceptable relative to the belief set G' and 'is a member of the belief set G' on the right-hand side of the Ramsey test for *Op* differ in extension (call this the *second route* out); for otherwise all the assumptions in Gärdenfors' theorem would be satisfied again, and a contradiction would be bound to follow. So among the original premises which led to Gärdenfors' contradictory conclusion, either Gärdenfors' version of the Ramsey test is dismissed because *Op* is not a nontrivial AGM belief revision operator, or it is dismissed because on its right-hand side '$A > B \in G$' is replaced by the nonequivalent 'is acceptable relative to the belief set G (and *Op*)'. In a nutshell, while the axioms of AGM belief revision may still be applied to such conditionals $A > B$ and to belief sets that have them as members, as these conditionals express propositions in the standard sense, one cannot assume Gärdenfors' version of the Ramsey test to hold, nor necessarily any of the Nontriviality assumptions.

2.3.1 The first route out If one takes the first route out, then there are several options of which Ramsey-test governed supposition operator *Op* to choose other than an AGM $*$, and since the conditionals in question are assumed to play role (ii), too, it is a natural option to consider 'is acceptable relative to the belief set G' to coincide both intensionally and extensionally with 'is a member of the belief set G' (although this is not entailed). For instance,

> $B \in G \star A$ iff $A > B$ is a member of the belief set G
>
> where $\star$ is the so-called belief update operator of Grahne [6] and Katsuno & Mendelzon [9],

is a viable and (as one can prove easily) perfectly consistent option; of course, the axioms of belief update in this sense differ from the standard axioms of belief revision. A very plausible way of understanding the differences between this form of belief update and AGM belief revision is in terms of the *type* of supposition that their corresponding Ramsey tests encode: while belief update corresponds to an *open-with-respect-to-the-facts* or perhaps *contrary-to-the-facts* supposition, AGM belief revision corresponds to a *matter-of-fact* supposition (see Joyce [8] for more on this distinction); in particular, if G is what the agent knows to be true in a strict sense of the word, and if A is consistent with G, then assuming A (say, 'Oswald did not kill Kennedy') as a matter of fact should lead to the expansion of G by A just as if A had been learned as a new piece of evidence about the actual world (though really it is hypothetical), while assuming A without committing oneself to assume that A is true in the actual world or perhaps even against what one thinks holds true in the actual world should not necessarily do so; this is exactly what happens in the cases of belief revision and belief update, respectively (which also corresponds to the different assessments of Ernest Adams' famous 'If Oswald did not kill Kennedy, someone else did' and 'If Oswald had not killed Kennedy, someone else would have'). While belief update is suitable for the counterfactual if-then sign—which is why '$>$' might be replaced by David Lewis's symbol for the subjunctive if-then that is governed by his semantics for counterfactuals—belief revision matches the indicative if-then sign; but we will not enter this topic here any further.

Another instance of the same route out of the problem is given by

> $B \in G + A$ iff $A > B$ is a member of the belief set G
>
> where $+$ is the standard AGM belief expansion operator.

In this case, '$>$' may simply be replaced by the material if-then sign, by the classical deduction theorem (for the expansion belief set $G + A$ is really just the deductive closure of $G \cup \{A\}$). Since expansion is actually not a bad model at all of *matter-of-fact* supposition, as long as one assumes that the antecedent of an indicative conditional is not plainly disbelieved by an agent—which is normally not the case at least for the indicative conditionals that are actually asserted by an agent—this goes some way to explaining why material conditionals are not so bad at all as logical representatives of indicative conditionals on a qualitative scale, as the classical "horseshoe analysis" of indicative conditionals had it. One can show that the triviality conclusion is avoided in this case by dropping one and only axiom of belief revision—the Consistency postulate 'For all G, A: if A is consistent then $G * A$ is consistent'—and maybe also by giving up at the same time on the counterparts of the original nontriviality requirements for belief revision (depending on how these Nontriviality Assumptions have been formulated): indeed, belief expansion is nothing but a degenerate version of belief revision in which either the supposition is consistent with what is believed by the agent or the result of the "revision" is simply the unique inconsistent belief set. Note that in either of the cases $Op = \star$ and $Op = +$, the conditionals $A > B$ on the right-hand sides of the corresponding Ramsey test variants can be members of belief sets under a supposition, and they may also be allowed to function as suppositions themselves by which belief sets may be changed, since these conditionals express propositions in the standard sense. As in Section 2.1, the relevant logical concepts for these conditionals have to be determined first, but of course this is simple here: as we said before, for $Op = \star$ the intended interpretation of the sign $>$ is the one of Lewis's subjunctive if-then, the logic of which is Lewis's preferred conditional logic VC, while for $Op = +$ the sign $>$ acquires the interpretation of the material implication sign, the logic of which is given by standard propositional logic.

We will not consider further possible choices for Op here. Let it suffice to say that if what we have said in this section so far is right, then we have the following result: If Op represents supposition open-with-respect-to-the-facts or contrary-to-the-facts, the axioms of belief revision do not hold for Op but the Ramsey test with belief instead of acceptability does so, and so will the nontriviality assumptions on Op (and $>$ is the subjunctive if-then). If Op represents the matter-of-fact supposition of statements A only that are consistent with what the agent believes, then again the axioms of belief revision do not hold for Op and maybe the nontriviality requirements will not do so either, but the Ramsey test with belief instead of acceptability does hold (and $>$ is the material if-then). Both of these options fall under the first route out of case 3 (as covered by the current Section 2.3). Finally, if Op represents the matter-of-fact supposition of arbitrary statements A, including such that are inconsistent with what the agent believes, then the axioms of belief revision do hold for Op as do the nontriviality constraints, the Ramsey test with belief instead of acceptability does not hold for Op, while the Ramsey test formulated in terms of acceptability does (and $>$ is then the general indicative if-then). This option either belongs to the purely suppositional ones covered by Section 2.2 on Ramsey-test conditionals, or it belongs to the second route out of case 3 to which we turn now.

2.3.2 The second route out The other option of maintaining consistency or nontriviality along the second route out is unexplored, as far as we know: it would require us not to assume about an agent, as we did in the first route that was just discussed, that her belief set at time t coincides with the agent's acceptability set at t. While the AGM postulates hold for Op and belief sets (remember that Op is a nontrivial AGM belief revision operator $*$ in the second route out), the Ramsey test concerns Op and acceptability sets according to this option; and since belief-to-be-true and pragmatic acceptability are assumed to come apart extensionally, no problematic consequences would have to follow from this bifurcation. Indeed, one would end up with exactly the second variant of the Ramsey test in Section 2.2, the only difference being that according to the current option one could also sensibly and truthfully *believe* conditionals such as $A > B$ *to be true*, it is just that the right-hand side of the Ramsey test does not make use of the availability of the concept of belief for conditionals of this kind, and it could not do so barring inconsistency or triviality. One worry about this second route out is that when an agent *asserts* a conditional of that sort, then the following two open questions arise. First, was the agent's belief in the truth of the conditional the necessary condition which had to be satisfied in order for the conditional to be assertable for the agent, or was it the acceptability of the conditional to the agent, or both? Secondly, does the speaker intend her audience to change their epistemic states such that they believe in the truth of the asserted conditional, or rather that the conditional is acceptable to them, or both? These different possibilities would have to be discussed in detail, and presumably the pragmatic context of assertion would have to be used as a cue as to which of these possibilities get realized.[8] A different worry concerns the logic of $>$ in this case: it could be determined either from the truth conditions for formulas of the form $A > B$ or from their acceptability conditions, and there would be no guarantee that the resulting logics would coincide.

Let us sum up what has been achieved in this section: we have determined in each out of three logically possible cases which of the premises in our revised arguments from the beginning of this section would have to be given up in order to restore consistency or nontriviality again. We will now turn in much more detail to the second case (as covered by Section 2.2 above): as we pointed out there, some of the relevant AGM postulates on Op and/or maybe even some of the nontriviality assumptions on Op must fail to hold for such Ramsey-test conditionals. But which ones exactly? As we will see, this depends on the exact formulation of the inconsistency or triviality result for $*$ that one intends to analyze (Gärdenfors' original one in his 1986 paper or someone else's). In the following, we will build a class of formal toy models in which Op will be given by a spheres system, which by Grove's theorem can be used to represent an AGM belief revision operator. However, we will formulate this operator from the start for acceptability sets rather than standard belief sets, where the acceptability of conditionals will be given again in terms of a formulation of the Ramsey test in which 'is acceptable relative to the belief set G (and Op)' replaces 'is a member of the belief set G'. In models of this class we will then determine which of the counterparts of standard AGM postulates and which of the nontriviality assumptions (if any) fail to hold; since the models are plausible and natural ones from the viewpoint of the case of Ramsey-test conditionals that we introduced in Section 2.2, we will be able to conclude that these axioms or nontriviality requirements are no longer satisfied by such conditionals and by the acceptability sets of which they are members. Hence we will finally be able to see clearly how Ramsey-test conditionals avoid inconsistency or triviality in spite of Gärdenfors' theorem.

3 A Formal Framework for Ramsey-test Conditionals

In the following, we assume some familiarity with the standard terminology that is used in theoretical accounts of belief revision. For the sake of simplicity, we restrict ourselves to a fixed propositional language $\mathcal{L}_{\text{prop}}$ which contains only finitely many propositional variables, and we let $\mathcal{L}_{\text{cond}}$ be the set of expressions $A > B$ where A and B are members of $\mathcal{L}_{\text{prop}}$. '$A$', '$B$', '$C$' always run over arbitrary sentences of $\mathcal{L}_{\text{prop}}$. We call the members of $\mathcal{L}_{\text{prop}}$ 'factual sentences' or 'factual formulas' and the members of $\mathcal{L}_{\text{cond}}$ simply 'conditionals', where factual formulas are assumed to be true or false at worlds and to express propositions, whereas conditionals are supposed to have a suppositional role to play but not a descriptive one, as explained in Section 2.2 above. In other words, our conditional formulas are meant to represent Ramsey-test conditionals. Indeed, we are going to carry out in detail now what has been said there about these conditionals for the case in which *Op* is a nontrivial AGM belief revision operator $*$. Finally, we choose $\mathcal{L}$ to be the union of $\mathcal{L}_{\text{prop}}$ and $\mathcal{L}_{\text{cond}}$. We are neither going to deal with applications of propositional connectives to conditionals nor with propositional compositions of conditionals and propositional formulas nor with nestings of the conditional operator. None of these restrictions is crucial as far as the derivability or nonderivability of the triviality theorems is concerned, but by using this simpler framework we may circumvent several notorious questions about the semantic status of complex sentences that involve conditionals, such as negated conditionals or conditionals of the form $(A > B) > C$. By restricting ourselves to finitely many propositional variables, we also do not have to deal with the complexities of infinite rankings of possible worlds.

A world of $\mathcal{L}_{\text{prop}}$ is just an assignment of truth values to the propositional variables in $\mathcal{L}_{\text{prop}}$; this assignment is then extended to yield an evaluation of all the formulas in $\mathcal{L}_{\text{prop}}$ at that world by the usual clauses of the semantics of propositional logic. Hence, for every world $w \in W$ and every formula $A \in \mathcal{L}_{\text{prop}}$, it will either be the case that $w \models A$ or that $w \models \neg A$.

By our assumptions on $\mathcal{L}_{\text{prop}}$, the set W of worlds of $\mathcal{L}_{\text{prop}}$ is finite. We consider a nontrivial belief state $\mathfrak{S}$ to be a tuple

$$\langle X_1, X_2, \ldots, X_n \rangle$$

where $\{X_1, X_2, \ldots, X_n\}$ is a partition of W; that is, (i) each X_i is a nonempty subset of W, (ii) $\bigcup_i X_i = W$, and (iii) if $i \neq j$ then $X_i \cap X_j = \varnothing$. Intuitively, $\mathfrak{S}$ is a linear preorder of worlds according to their degree of plausibility: the worlds in X_1 are more plausible than the worlds in X_2, those in X_2 are more plausible than the members of X_3, and so forth. The worlds within each "layer" X_i are supposed to be equally plausible. If a world w is a member of X_i, we say that $rk_{\mathfrak{S}}(w) = i$ ("w has rank i in $\mathfrak{S}$") and, accordingly, that layer X_i is of rank i. For $A \in \mathcal{L}_{\text{prop}}$, $[A]$ is the set of worlds in W which satisfy A; the worlds in $[A]$ are sometimes called A-worlds or $[A]$-worlds. Just as in the possible worlds semantics for doxastic logic, sets of worlds are to be regarded as the semantic counterparts of belief contents and indeed for most purposes $[A]$ and A can be "identified" with each other. If X is a set of worlds, let $\min_{\mathfrak{S}}(X)$ be the set of worlds in X which have minimal rank in $\mathfrak{S}$ among the members of X; that is, $\min_{\mathfrak{S}}(X) = \{w \in X \mid \nexists w' \in X : rk_{\mathfrak{S}}(w') < rk_{\mathfrak{S}}(w)\}$.

For technical reasons, we also introduce a unique trivial belief state $\mathfrak{S}_{\perp}$ which is not to be identical with any of the nontrivial belief states. Intuitively, $\mathfrak{S}_{\perp}$ is the "inconsistent" belief state in which everything whatsoever is believed or accepted. The total set of belief states is the union of the set of nontrivial belief states with $\{\mathfrak{S}_{\perp}\}$.

A nontrivial belief state $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ contains two sorts of epistemic information: belief, which is represented by the minimal layer X_1, and acceptability, which is given by the ranking of the worlds in terms of layers. More formally, we define it as follows.

Definition 3.1 For all nontrivial belief states $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$,

Factual Belief For all sentences A in $\mathcal{L}_{\text{prop}}$,

$$\mathfrak{S} \models A \text{ ("}A\text{ is believed true in }\mathfrak{S}\text{") :iff } X_1 \subseteq [A].$$

Acceptability For all conditionals $A > B$ in $\mathcal{L}_{\text{cond}}$,

$$\mathfrak{S} \models A > B \text{ ("}A > B\text{ is acceptable relative to }\mathfrak{S}\text{") :iff } \min_{\mathfrak{S}}([A]) \subseteq [B].$$

Belief in a sentence A which expresses a proposition is thus a belief of what is most plausibly the case: the most plausible worlds in the given belief state satisfy A. The propositional content of that belief is the set of worlds in which A is true, that is, $[A]$. The acceptability of a conditional amounts to regarding the most plausible worlds among those that satisfy A to satisfy B. As intended, these conditionals $A > B$ do not express propositions, which shows up in the way that they are not evaluated at worlds and that therefore '$[A > B]$' is undefined.

In the case of the trivial belief state, we simply assume $\mathfrak{S}_{\perp} \models A$ and $\mathfrak{S}_{\perp} \models A > B$ for all $A \in \mathcal{L}_{\text{prop}}$, $A > B \in \mathcal{L}_{\text{cond}}$.

Variations of the formal models of belief states that we use—sometimes called "spheres systems" or "ranked models"—have been employed successfully in various formal accounts of knowledge representation, belief revision, and nonmonotonic reasoning (Grove [7], Spohn [23], and Kraus et al. [10] are typical references). The ranks of worlds can be interpreted differently: the higher the rank of a world, the less (i) plausible, (ii) normal, (iii) entrenched, or (iv) expected it is supposed to be. Ranking worlds may also be viewed as assigning probabilistic orders of magnitude to them (cf. Lehmann & Magidor [11]). As we mentioned before, Grove [7] proved that AGM belief revision operators $*$ and spheres systems as the above one stand in a one-to-one correspondence to each other, where '$B \in G * A$' is represented as in our clause for the acceptability of $A > B$ above, and where the corresponding belief set (*not* acceptability set) G is represented as in our clause for factual belief above.

This formal account of belief states could be extended and made more fine-grained in several respects: we could allow the set of worlds in a belief state to be a proper subset of W; the restrictions on plausibility preorders for worlds could be weakened such that a state would be a preferential model based on an arbitrary strict partial order (cf. Kraus et al. [10]) rather than a linear ranked model; and so forth. But since these additional degrees of complexity turn out to be unnecessary for our analysis of the triviality results, we avoid the complications.

If $\mathfrak{S}$ is a belief state, let

$$\text{Th}(\mathfrak{S}) = \{A \in \mathcal{L}_{\text{prop}} : \mathfrak{S} \models A\} \cup \{A > B \in \mathcal{L}_{\text{cond}} : \mathfrak{S} \models A > B\}$$

be the acceptability set that is associated with or determined by $\mathfrak{S}$. Hence, as anticipated in Section 2.2, what is believed to be true by an agent in a state is included in what is accepted by the agent in that state. In the following, when we speak of acceptability sets in a semantic context, we always refer to subsets G of $\mathcal{L}$ for which there is a belief state $\mathfrak{S}$ such that $G = \text{Th}(\mathfrak{S})$; '$G$', '$H$', '$K$' are going to denote arbitrary acceptability sets. Note that we use 'G' now to refer to acceptability sets where we used it mostly to refer to *belief* sets in the first two sections of this paper. Let us state this characterization of acceptability sets explicitly.

Definition 3.2 (Acceptability Sets) G is an acceptability set if and only if there is a belief state $\mathfrak{S}$ such that

$$G = \text{Th}(\mathfrak{S}).$$

It is easy to see that if $\mathfrak{S}_1, \mathfrak{S}_2$ are belief states which are distinct, that is, $\mathfrak{S}_1 \neq \mathfrak{S}_2$, then also $\text{Th}(\mathfrak{S}_1)$ differs from $\text{Th}(\mathfrak{S}_2)$. In this sense, belief states and their corresponding acceptability sets may be "identified" with each other as well. Accordingly, if $G = \text{Th}(\mathfrak{S})$, then instead of saying that A is believed in $\mathfrak{S}$ we may just as well say that A is believed in G, which in turn is equivalent to saying that $A \in G$; analogously for conditionals $A > B$ and acceptability instead of belief. Note that for all factual sentences A and for all acceptability sets G it follows that $A \in G$ if and only if $\top > A \in G$ where '$\top$' denotes the logical verum, that is, a fixed propositional logical truth. Moreover, for all factual sentences A and for all acceptability sets G, the conditional $\bot > A$ is in G, where '$\bot$' denotes the logical falsum, that is, a fixed propositional contradiction. Finally, all acceptability sets are deductively closed with respect to factual formulas, and the set of factual members of an acceptability set is consistent if and only if the acceptability set is associated with a nontrivial belief state. $\text{Th}(\mathfrak{S}_\bot)$ will be denoted by '$G_\bot$'; the set of factual members of $G_\bot$ is inconsistent in the standard propositional sense; in fact, $G_\bot = \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{cond}}$. $G_\bot$ is called the 'trivial acceptability set', all other acceptability sets are called 'nontrivial', and the total set of acceptability sets is simply the set of the trivial acceptability set and the nontrivial beliefs set taken together. Our assumption that for every possible world w for $\mathcal{L}_{\text{prop}}$ and for every belief state $\mathfrak{S}$ there is some layer of $\mathfrak{S}$ that contains w as a member has the consequence that a conditional of the form $A > \bot$ can only be a member of a nontrivial belief state if A is logically false. In this sense, our belief states are "logically cautious": a logical possibility is never totally excluded in a belief state although it may of course be viewed as highly implausible. (This is in line with the AGM axioms, if spheres systems are replaced by the belief revision operators that they represent.)

As we outline in the Appendix, one can define syntactic notions of consistency and derivability as well as semantic notions such as satisfiability and logical consequence for arbitrary subsets and formulas of our full language $\mathcal{L}$ such that (i) all acceptability sets contain all formulas—whether factual or conditional—that are derivable from them, and (ii) an acceptability set is consistent according to this notion of consistency if and only if it is different from $G_\bot$, that is, if it is nontrivial. The logic is Ernest Adams' logic for indicative conditionals again. Moreover, one can prove that an acceptability set G is consistent with a factual formula A if and only if the set of factual members of G is consistent with A. In other words, the rules of Adams' logic are precisely the ones that are sound and complete with respect to models as discussed in this section, in the sense that a formula is derivable by these rules from a finite set of formulas if and only if for every model as above, if the model satisfies all the formulas in the set, then it also satisfies the derived formula. (This formal result is, of course, well known from Adams' own work in Adams [1], but also from Kraus et al. [10] and other sources.) In Sections 4 and 5, we usually prefer not to speak of the consistency of an acceptability set but rather of its being distinct from $G_\perp$ and we do not say that an acceptability set G is consistent with a factual formula but rather that the set of factual members of G is consistent with the factual formula. In this way, 'consistent' may always be understood in the standard sense of propositional logic.

4 Minimal Requirements for Revision

Now we turn to principles of belief revision for formulas in $\mathcal{L}_{\text{prop}}$ and acceptability sets. Each such principle will be stated in a syntactic version and a semantic version, where the semantic versions have a double function: constraints are often more transparent and more easily justified if stated semantically; at the same time, the syntactic propositions can be seen to be consistent with all previously introduced principles by considering their semantic counterparts. While the syntactic principles are not tied to any semantic construction at all, their semantic counterparts will be tied directly to the formal framework that was developed in Section 3. In all the semantic principles, the term 'acceptability set' receives its meaning through the semantic construction of the last section. In the syntactic principles, one may instead consider an acceptability set to be a set of formulas in $\mathcal{L}$ that is closed under the rules of Adams' logic of indicative conditionals as explained in the Appendix (where maybe an acceptability set has to satisfy some additional constraints over and above closure under Adams' rules). The postulates presuppose some underlying set of acceptability sets in that sense.

As intended in Section 2.2, the revision mapping is regarded as a function of the following kind.

Revision Operator: Syntactic Version, Part I

> If $\mathcal{B}$ is the set of acceptability sets, then $*$ is a mapping of the form
> $* : \mathcal{B} \times \mathcal{L}_{\text{prop}} \rightarrow \mathcal{B}$.

Furthermore, we want $*$ to depend only on the semantic aspects of those sentences by which an acceptability set is to be revised. Therefore, we add

Revision Operator: Syntactic Version, Part II

> For all acceptability sets G, for all factual sentences A, B:
> if A is logically equivalent with B, then $G * A = G * B$.

Since acceptability sets correspond to belief states and factual sentences correspond to sets of worlds, that is, to members of the power set $\wp(W)$ of W, this last constraint allows us to rewrite Part I of our syntactic assumption on the revision operator in semantic terms.

Revision Operator: Semantic Version

> If $\mathfrak{B}$ is the set of belief states, then $*$ is a mapping of the form
> $* : \mathfrak{B} \times \wp(W) \rightarrow \mathfrak{B}$.

At the same time, this covers semantically Part 2 from above, too, since logically equivalent formulas express the same proposition. While we speak semantically of the revision of a belief state by a formula, this should not be mixed up with the revision of a belief set by a formula on the syntactic level: syntactically, it is *acceptability* sets which get revised.

These assumptions on revision operators do not achieve much more than determining the set-theoretic status of these mappings. An obvious way of strengthening them is to postulate that the revision by a consistent formula can only lead to a nontrivial acceptability set, even if the acceptability set to which the revision operator is applied is itself the trivial one. This yields

Consistency: Syntactic Version

For all factual sentences A, for all acceptability sets G:
If A is consistent, then $G * A$ is nontrivial; that is, $G * A \neq G_\perp$ (and correspondingly the set of factual members of $G * A$ is consistent[9]).

The corresponding semantic principle is given by turning from consistent sentences A and acceptability sets G on the one hand to nonempty sets $Y = [A]$ and belief states $\mathfrak{S}$ with $\mathrm{Th}(\mathfrak{S}) = G$ on the other.

Consistency: Semantic Version

For all $Y \subseteq W$, for all belief states $\mathfrak{S}$:
If Y is nonempty, then $\mathfrak{S} * Y$ is nontrivial; that is, $\mathfrak{S} * Y \neq \mathfrak{S}_\perp$.

Revision by *inconsistent* information is a singular and somewhat less interesting case which is not so relevant as far as the aims of this paper are concerned. For the sake of completeness and since it is convenient as far as the presentation of some of our later principles are concerned, we add

Inconsistency: Syntactic Version

For all factual sentences A, for all acceptability sets G:
If A is inconsistent, then $G * A$ is trivial; that is, $G * A = G_\perp$ (and correspondingly the set of factual members of $G * A$ is inconsistent).

Inconsistency: Semantic Version

For all $Y \subseteq W$, for all belief states $\mathfrak{S}$:
If Y is empty, then $\mathfrak{S} * Y$ is trivial; that is, $\mathfrak{S} * Y = \mathfrak{S}_\perp$.

This weak account of belief revision can be extended by introducing a constraint which we call the 'Deconditionalization principle' and which we state first in its semantic version.

Deconditionalization: Semantic Version

For all nonempty $Y \subseteq W$, for all belief states $\mathfrak{S}$:

(i) If $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ is nontrivial, then the set of most plausible worlds in $\mathfrak{S} * Y$ is identical to $\min_{\mathfrak{S}}(Y)$, that is, the set of most plausible worlds in $\mathfrak{S}$ among those worlds that are included in Y.

(ii) If $\mathfrak{S}$ is trivial, then the set of most plausible worlds in $\mathfrak{S} * Y$ is some subset of Y.

We will postpone our comment on (ii) and our syntactic statement of Deconditionalization. With regards to (i), if the rank of the members of $\min_{\mathfrak{S}}(Y)$ is k, then the minimal layer in $\mathfrak{S} * Y$ is demanded to be $X_k \cap Y$. Note that if Y is nonempty, $\min_{\mathfrak{S}}(Y)$ is nonempty, too, because every possible world for $\mathcal{L}_{\mathrm{prop}}$ is a member of some layer in $\mathfrak{S}$ by the definition of belief states.

In the case of revising a nontrivial belief state, Deconditionalization leaves open how worlds are ranked in $\mathfrak{S} * Y$ except for the single constraint that the worlds that have been regarded as most plausible among all the Y-worlds before the revision should turn out to be the most plausible worlds *simpliciter* after the revision. If Deconditionalization were not accepted, it would be questionable what it meant to say in the first place that $\min_{\mathfrak{S}}(Y)$ was the set of worlds which were *most plausible among all the Y-worlds*. What else could the most plausible worlds in $\mathfrak{S} * Y$ be like? If not all of them were worlds that satisfied A where $Y = [A]$, then the following consequence of our syntactic version of Deconditionalization to come (given the Inconsistency principle and the derivability of $A > A$ in Adams' logic) would be invalidated.

Success: Syntactic Version

For all factual sentences A, for all acceptability sets G:
$A \in G * A$.

Success simply states that the revision of a belief state by a sentence A leads to a belief state in which A is believed. The main function of part (ii) of Deconditionalization is to ensure that $A \in G * A$ follows even in cases where G is indeed trivial.

Success might be viewed as an oversimplification of real epistemic situations, in which it is definitely not always the case that new evidence is automatically regarded as "most plausible." But as we said before, revision is here really revision *by supposition* rather than by learning evidence, which is why Success is a given. The Success condition has the following semantic version, which follows from the semantic version of Deconditionalization (and Inconsistency).

Success: Semantic Version

For all subsets Y of W, for all belief states $\mathfrak{S}$:
If $\mathfrak{S} * Y$ is nontrivial, then Y is a superset of the minimal layer of $\mathfrak{S} * Y$.

Given that Success is to be valid, all of the most plausible worlds in $\mathfrak{S} * Y$ have to be Y-worlds.

Now let us continue our analysis of Deconditionalization: assume all of the most plausible worlds in $\mathfrak{S} * Y$ are worlds that are members of $Y = [A]$, since we want Success to hold. But, say, some of these worlds would not be members of the set of *minimal* A-worlds, $\min_{\mathfrak{S}}[A]$: why would these worlds all of a sudden be regarded as equally plausible or even more plausible than the worlds in $\min_{\mathfrak{S}}[A]$? Accordingly, if not all the worlds in $\min_{\mathfrak{S}}[A]$ were among the most plausible worlds in $\mathfrak{S} * Y$ but only some of the worlds in $\min_{\mathfrak{S}}[A]$ were, why would the former members of $\min_{\mathfrak{S}}[A]$ be no longer considered equally plausible? It seems that if the plausibility interpretation of the ranking of worlds in belief states is to be taken seriously, the Deconditionalization principle is inevitable and so are its logical consequences.

This is where the Ramsey test comes into play. The Ramsey test is simply a part of the syntactic version of Deconditionalization which says

Deconditionalization: Syntactic Version

For all consistent factual sentences A, for all acceptability sets G:

(i) (Ramsey test)
if G is nontrivial, then for all factual sentences B:
$B \in G * A$ if and only if $A > B \in G$;

(ii) if G is trivial, then $A \in G * A$.

Usually, regarding (i), the Ramsey test is stated without restricting 'A' to *consistent* sentences and without restricting 'G' to *nontrivial* acceptability sets. Although the restrictions that we have introduced are irrelevant for our later discussion of the inconsistency results, they are necessary given the way in which we have treated trivial acceptability sets and revision by inconsistent information, since in Adams' logic any conditional of the form $\bot > B$ is derivable in every acceptability set, and the trivial (that is, inconsistent) acceptability set includes all conditionals whatsoever; see the Appendix.

Note that we could have expressed clause (i) of the semantic version of Deconditionalization just as well in the following form which is closer to the syntactic structure of the Ramsey test: If $\mathfrak{S}$ is nontrivial and $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$, then $\mathfrak{S} \models A > B$ if and only if $\mathfrak{S} * A \models B$.

What does the Ramsey test express as far as $*$ is concerned? It says that if $\mathfrak{S}$ is a nontrivial belief state and if the most plausible worlds in $\mathfrak{S}$ among those worlds that satisfy A also satisfy B, then the worlds in $\mathfrak{S} * [A]$ that are most plausible *simpliciter* ought to satisfy B, and vice versa. But this is precisely what the Deconditionalization principle says which demands that the set of worlds that are most plausible in $\mathfrak{S} * [A]$ is identical to the set of worlds that are most plausible in $\mathfrak{S}$ among all the A-worlds. Given the intended interpretation of our formal framework and given the rest of our minimal restrictions on revision, the Deconditionalization principle and thus the Ramsey test are well justified and do not lead to any problems or even to inconsistency.

As Gärdenfors has observed, the Ramsey test (together with Inconsistency) implies that $*$ satisfies the following interesting monotonicity principle.

Monotonicity: Syntactic Version

> For all factual sentences A, for all acceptability sets G, G' where G' is nontrivial:
> If $G \subseteq G'$, then the set of factual members of $G * A$ is a subset of the set of factual members of $G' * A$.

Gärdenfors does not restrict G' to nontrivial acceptability sets, which is again not relevant for the topic of the Ramsey test. More importantly, he states Monotonicity in the simpler and much stronger form 'If $G \subseteq G'$, then $G * A \subseteq G' * A$', which seems to indicate that *all* members of $G * A$ are also members of $G' * A$, including all the conditionals in $G * A$. However, this would only follow from the Ramsey test if nestings of conditionals were allowed to occur as members of acceptability sets. Since we have restricted the antecedents and consequents of conditionals to factual formulas, the Ramsey test only implies the kind of Monotonicity principle that we have stated above. But note that Gärdenfors himself only needs our version of the Monotonicity principle in the proof of his triviality theorem and the same holds for Lindström and Rabinowicz' [16] exposition of the same result. As Gärdenfors correctly remarks on the same page (cf. Gärdenfors [5, p. 157]), "it is not necessary to assume that $\mathbf{L}'$ [the object language in question] contains sentences with iterated occurrences of the conditional operator."

In order to state the semantic version of Monotonicity in a transparent manner, it is useful to introduce some additional terminology. Let us define $\ll$ to be a binary relation on belief states such that $\mathfrak{S} \ll \mathfrak{S}'$ ("$\mathfrak{S}'$ is a refinement of $\mathfrak{S}$") if and only if for all $w, w' \in W$: if $rk_{\mathfrak{S}}(w) < rk_{\mathfrak{S}}(w')$ then $rk_{\mathfrak{S}'}(w) < rk_{\mathfrak{S}'}(w')$. Refinements are not necessarily meant to be strict; instead of saying that $\mathfrak{S}'$ is a refinement of $\mathfrak{S}$, we also say that $\mathfrak{S}$ is equally or more coarse-grained than $\mathfrak{S}'$. If $\mathfrak{S} \ll \mathfrak{S}'$, then $\mathfrak{S}'$ has to be the result of preserving all the ordering relationships between worlds in $\mathfrak{S}$ while possibly adding new ones by dividing some of the X_i sets of $\mathfrak{S}$ into several layers. One can show that every refinement of a belief state corresponds to an increase of its associated acceptability set and vice versa; that is, for all belief states $\mathfrak{S}, \mathfrak{S}'$ holds: $\mathfrak{S} \ll \mathfrak{S}'$ if and only if $\mathrm{Th}(\mathfrak{S}) \subseteq \mathrm{Th}(\mathfrak{S}')$. Hence, the syntactic version of Monotonicity has a semantic counterpart that can be expressed in terms of $\ll$.

Monotonicity: Semantic Version

For all $Y \subseteq W$, for all belief states $\mathfrak{S}, \mathfrak{S}'$, where $\mathfrak{S}'$ is nontrivial;
let $\mathfrak{S} * Y = \langle X_1, X_2, \ldots, X_m \rangle$, $\mathfrak{S}' * Y = \langle X'_1, X'_2, \ldots, X'_n \rangle$:
if $\mathfrak{S} \ll \mathfrak{S}'$, then $X_1 \supseteq X'_1$.

The antecedent condition of Monotonicity is actually a strong constraint on $\mathfrak{S}, \mathfrak{S}'$, since if $\mathfrak{S}'$ is assumed to be a refinement of $\mathfrak{S}$, then all order relations between worlds in $\mathfrak{S}$ also have to hold in $\mathfrak{S}$.

If we allowed for belief states $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ which would exclude certain possible worlds, that is, where $\bigcup_i X_i \subsetneq W$, then clause (i) of the semantic Deconditionalization principle would have to be changed into 'If $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ is nontrivial *and* $Y \cap \bigcup_i X_i \neq \varnothing$, then the set of most plausible worlds in $\mathfrak{S} * Y$ is identical to the set of most plausible worlds in $\mathfrak{S}$ among those worlds that are included in Y'. The corresponding formulation of the Ramsey test would be 'If G is nontrivial *and* $A > \bot \notin G$, then for all factual sentences B: $B \in G * A$ if and only if $A > B \in G$'. Moreover, a weaker Monotonicity principle would be implied which we state in its syntactic version: For all factual sentences A, for all acceptability sets G, G' where G' is nontrivial *and* $A > \bot \notin G'$: If $G \subseteq G'$, then the set of factual members of $G * A$ is a subset of the set of factual members of $G' * A$.

In the next section we will extend our assumptions on $*$ by more controversial principles which nevertheless conform in some sense to the conception of revision as belief change subject to minimal mutilation.

5 A Possible Way of Strengthening the Revision Postulates

All of our considerations on $*$ in this section will solely deal with the revision of a belief state $\mathfrak{S}$ by a consistent sentence A such that A is also consistent with the set of factual members of $\mathrm{Th}(\mathfrak{S})$.

Semantically, the consistency of A with the set of factual members of $\mathrm{Th}(\mathfrak{S})$ corresponds to the fact that the layer of $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ which has minimal rank, that is, X_1, contains at least one A-world as a member, or, equivalently, $X_1 \cap [A] \neq \varnothing$. What should the revision of $\mathfrak{S}$ by $Y = [A]$ look like in such a situation? Here is what we regard as one possible plausible answer to this question.[10]

Minimal Mutilation: Semantic Version

For all $Y \subseteq W$, for all nontrivial belief states $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$ such that $X_1 \cap Y \neq \varnothing$:

(i) If $X_1 \cap (W \setminus Y) \neq \varnothing$, then $\mathfrak{S} * Y$ is of the form

$$\langle X_1 \cap Y, X_1 \cap (W \setminus Y), X_2, \ldots, X_n \rangle.$$

(ii) If $X_1 \cap (W \setminus Y) = \varnothing$, that is, $X_1 \subseteq Y$, then $\mathfrak{S} * Y$ is identical to $\mathfrak{S}$.

Let us interpret this principle: before the revision takes place, the worlds in X_1 are considered more plausible according to $\mathfrak{S}$ than all other worlds. Then $\mathfrak{S}$ is revised by the proposition Y: this confirms the $X_1 \cap Y$-worlds having minimal rank among all worlds; such worlds exist by the antecedent assumption of the Minimal Mutilation principle. By the Success principle, if there are further worlds in X_1, that is, if there are worlds in $X_1 \cap (W \setminus Y)$, then they have to be removed from the first layer, so they should be placed on a layer that is of higher rank than $X_1 \cap Y$; but since they were also regarded as more plausible than all the worlds in $X_2 \cup \cdots \cup X_n$, they should be in a layer that is of lesser rank than $X_2 \cup \cdots \cup X_n$. On the other hand, if there are no worlds in $X_1 \cap (W \setminus Y)$, then nothing has to be changed by revision. This is precisely what is postulated by the Minimal Mutilation principle.

Of course, there are some further possible options: for example, one might not use all worlds in $X_1 \cap (W \setminus Y)$ as the new second layer, but instead distribute $X_1 \cap (W \setminus Y)$ over several layers that are located between $X_1 \cap Y$ and X_2. But why should worlds that are indifferent with respect to both their former rank of plausibility and their logical exclusion of Y be treated differently after the revision by Y? A more serious objection is that although the worlds in $X_2 \cup \cdots \cup X_n$ have been regarded less plausible than the worlds in $X_1 \cap (W \setminus Y)$, some of them might actually be Y-worlds in contrast to the latter. Therefore, since Y is supposed to be "most plausible" after the revision, the plausibility of the worlds in $(X_2 \cup \cdots \cup X_n) \cap Y$ might actually be considered to rise in plausibility if compared to the worlds in $X_1 \cap (W \setminus Y)$ and perhaps also in comparison with the worlds in $(X_2 \cup \cdots \cup X_n) \cap (W \setminus Y)$. For example, a revision procedure might be applied by which all Y-worlds are ranked below all $W \setminus Y$-worlds. This is roughly what Spohn [23] demands when he makes use not just of the ordinal but also of the arithmetic properties of world ranks. Although this is an interesting and elegant suggestion, we hold on to a model in which plausibility ranks override the new incoming information except for the worlds in the first layer. This conforms to the assumption that Y is taken as a confirmation that the actual world is among the most plausible Y-worlds but that it neither confirms nor disconfirms any plausibility judgment about "abnormal" Y-worlds. Below we will add a further argument in favor of the Minimal Mutilation principle which proceeds from conservativeness or minimal mutilation and which also gives the principle its name.

There is a more or less direct translation of the semantic version of the Minimal Mutilation principle into syntactic terms; however, since the formulation is rather clumsy, we refrain from stating it but instead work our way toward another syntactic characterization of $*$ that turns out to be equivalent to the immediate syntactic version of Minimal Mutilation.

For now we want to see what the logical consequences of the semantic version of the Minimal Mutilation principle look like. The following is perhaps its most important implication.

Preservation: Syntactic Version

For all factual sentences A, for all nontrivial acceptability sets G:
If A is consistent with the set of factual members of G, then $G \subseteq G * A$.

Preservation expresses that the revision of an acceptability set by a factual sentence that is consistent with the acceptability set can only increase what the agent commits herself to but it can never decrease it. As Gärdenfors [5] outlines, Preservation in the purely propositional case also follows from the Bayesian account of belief change. Both Gärdenfors and Lindström & Rabinowicz [16] list Preservation among their prima facie acceptable postulates. The latter discuss different ways of getting around the inconsistency result, one of which is to abandon the Preservation principle at least in its full strength. Gärdenfors presents his own result in terms of a dilemma: either the Ramsey test or Preservation has to be given up. While he argues in favor of keeping the Preservation principle, our treatment of revision shows that neither has to be the case, once the path toward Ramsey-test conditionals in the sense of Section 2.2 has been taken. The validity of Preservation is another reason why we have characterized revision in terms of our Minimal Mutilation principle rather than in terms of, say, Spohn's suggestion, according to which Preservation does *not* hold. Since Preservation conforms to a rationale of informational economy itself—see Gärdenfors' discussion of this point in [5, p. 67 and p. 157]—its derivability from Minimal Mutilation also strengthens the conservativeness characteristic that we have ascribed to the latter. In this sense, by accepting Minimal Mutilation and therefore also Preservation, we are able to stick as closely as possible to some of the original intentions that were driving the theory of belief revision. This said, these intentions are certainly not unproblematic themselves (cf. Rott [19]), and AGM conceived of them as applying primarily to statements that express propositions, unlike the conditionals that we are working with at present.

The semantic version of Preservation gives some insight in why every revision operator that satisfies Minimal Mutilation must also satisfy the Preservation condition. We make use of our binary relation $\ll$ on belief states again.

Preservation: Semantic Version

> For all subsets Y of W, for all nontrivial belief states $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$:
> If $X_1 \cap Y \neq \varnothing$, then $\mathfrak{S} \ll \mathfrak{S} * Y$.

From what we have demanded in the Minimal Mutilation principle, it follows immediately that if $X_1 \cap Y \neq \varnothing$, then $\mathfrak{S} * Y$ is actually a refinement of $\mathfrak{S}$.

Our final observation on the revision of acceptability sets by factual sentences that are consistent with the acceptability set is about the "cautiousness" of revision; it is again implied by the Minimal Mutilation principle.

Cautiousness: Semantic Version

> For all subsets Y of W, for all nontrivial belief states $\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$:
> If $X_1 \cap Y \neq \varnothing$, then $\mathfrak{S} * Y$ is the most coarse-grained belief state that satisfies Success and Preservation.

Since Cautiousness is actually a unique characterization of the revision operator as far as inputs are concerned that are consistent with all current factual beliefs, its syntactic version is at the same time a nice substitute for the immediate syntactic version of Minimal Mutilation.

Minimal Mutilation & Cautiousness: Syntactic Version

> For all factual formulas A, for all nontrivial acceptability sets G:
> If A is consistent with the set of factual members of G, then $G * A$ is the least acceptability set that satisfies Success and Preservation.

This consequence again supports our interpretation of the Minimal Mutilation principle in terms of conservativeness: considered semantically, the ranking of worlds after a revision is demanded to be as similar as possible to the ranking before the revision, where the 'as possible' refers to the fact that Success is to be satisfied and where 'similar' corresponds to Preservation. In syntactic terms, what the agent holds after the revision should be as similar as possible to what she held before the revision, where the 'as possible' again refers to the fact that Success is to be satisfied and where 'similar' again corresponds to Preservation.

All of the above principles for belief revision are consistent with each other, as can be seen easily by considering their semantic versions. In fact, they are still weak, presumably *too weak*; in particular, the revision of acceptability sets G by sentences that are inconsistent with the factual members of G is strongly underdetermined. But the current list of principles is already sufficient for showing that some of the additional postulates that have been used in some of the well-known proofs of the inconsistency or triviality result on belief revision and the Ramsey test are false or at least highly questionable *for Ramsey-test conditionals*.

6 Lindström & Rabinowicz's Version of the Inconsistency Theorem

The presentation of Gärdenfors' inconsistency result in this section follows the lines of Lindström and Rabinowicz [16]. The premises that they use in order to derive the inconsistency theorem are as follows (using our own terminology wherever appropriate).

Ramsey Test If A is consistent and G is nontrivial, then
$A > B \in G$ if and only if $B \in G * A$.

Success $A \in G * A$.

Consistency If both G and A, when considered separately, are consistent, then $G * A$ is consistent.

Preservation If A is consistent with the set of factual members of G, then $G \subseteq G * A$.

Nontriviality There are two sentences B and C and three nontrivial acceptability sets G, H, and K such that

(i) $B \in G$ and the set of factual members of G is consistent with $\{\neg C\}$,
(ii) $C \in H$ and the set of factual members of H is consistent with $\{\neg B\}$,
(iii) $G \subseteq K$ and $H \subseteq K$.

As Lindström and Rabinowicz show, a contradiction can be derived from these assumptions taken together. We have added the restriction to nontrivial acceptability sets to their original version of the Ramsey test, which is unproblematic as they say themselves in Lindström and Rabinowicz [16], footnote 10. Furthermore, we demand A to be consistent. Preservation is actually stated by Lindström and Rabinowicz in the way that if A is consistent with G, then $G \subseteq G * A$. Since they do not uniquely determine their notion of consistency of an acceptability set with a factual formula, we have put Preservation in a form in which only the consistency of a set of propositional formulas with another propositional formula is needed; thus we rather postulate that if A is consistent with *the set of factual members of* G, then $G \subseteq G*A$. The consistency of an acceptability set is again to be identified with its being distinct from the trivial acceptability set, as can in fact be derived by presupposing Adams' logic as stated in the Appendix.

Given the framework and the principles that we have introduced in the last three sections, the Ramsey test, Success, Consistency, and Preservation are valid and consistent with each other. While Lindström and Rabinowicz [16] discuss reasons for giving up or weakening either Preservation or the Ramsey test and while they present various ways of to how achieve this, they do not question Nontriviality. In contrast, we claim that it is precisely the Nontriviality premise that is implausible *if one takes up our case 2 (Section 2.2) on Ramsey-test conditionals*. As we are going to show, for all belief states $\mathfrak{S}_1, \mathfrak{S}_2, \mathfrak{S}_3$ and formulas B, C in $\mathcal{L}_{\text{prop}}$, *Nontriviality is not satisfied* for $G = \text{Th}(\mathfrak{S}_1)$, $H = \text{Th}(\mathfrak{S}_2)$, $K = \text{Th}(\mathfrak{S}_3)$ and B, C. This result is independent of *any* assumptions on the revision operator; only some features of our formal background framework are presupposed. While this observation is interesting in itself, it certainly does not show that there is anything "trivial" about the arbitrarily chosen belief states $\mathfrak{S}_1, \mathfrak{S}_2, \mathfrak{S}_3$ and formulas B, C, despite the fact that the so-called Nontriviality condition is *not* satisfied. Since the notions of belief state and acceptability set that we have introduced above constitute reasonable formal representations of our corresponding pretheoretic notions, we conclude that Nontriviality is implausible, that it does *not* express that the epistemic setting we are dealing with is "nontrivialized", and thus that Lindström and Rabinowicz' version of the triviality result does not speak against the Ramsey test, once conditionals are understood to be purely suppositional (role (i), not role (ii)).

Consider three belief states,

$$\mathfrak{S}_1 = \langle X_1, X_2, \ldots, X_{n_1} \rangle, \mathfrak{S}_2 = \langle Y_1, Y_2, \ldots, Y_{n_2} \rangle, \mathfrak{S}_3 = \langle Z_1, Z_2, \ldots, Z_{n_3} \rangle,$$

such that the part of Nontriviality is instantiated which demands that $G = \text{Th}(\mathfrak{S}_1) \subseteq K = \text{Th}(\mathfrak{S}_3)$ and $H = \text{Th}(\mathfrak{S}_2) \subseteq K = \text{Th}(\mathfrak{S}_3)$. $G \subseteq K$ (analogously for $H \subseteq K$) really has two aspects, (i) for all $A \in \mathcal{L}_{\text{prop}}$: if $A \in G$, then $A \in K$, (ii) for all $A > B \in \mathcal{L}_{\text{cond}}$: if $A > B \in G$, then $A > B \in K$; the second aspect is the more important one. Since the factual subset of $G = \text{Th}(\mathfrak{S}_1)$ is consistent with $\neg C$ and since $B \in G$, it follows that $\mathfrak{S}_1 \models (\neg B \vee \neg C) > B$. Accordingly, since the factual subset of $H = \text{Th}(\mathfrak{S}_2)$ is consistent with $\neg B$ and since $C \in H$, it must be the case that $\mathfrak{S}_2 \models (\neg B \vee \neg C) > C$. Thus, because of $G, H \subseteq K = \text{Th}(\mathfrak{S}_3)$, we also have $\mathfrak{S}_3 \models (\neg B \vee \neg C) > B$ and $\mathfrak{S}_3 \models (\neg B \vee \neg C) > C$ and, therefore, $\mathfrak{S}_3 \models (\neg B \vee \neg C) > (B \wedge C)$. But at the same time $\mathfrak{S}_3 \models (\neg B \vee \neg C) > (\neg B \vee \neg C)$ must hold (trivially), which implies $\mathfrak{S}_3 \models (\neg B \vee \neg C) > \bot$. So there cannot be any worlds in W satisfying $\neg B \vee \neg C$. But W is the set of *all* possible worlds for $\mathcal{L}$; that is, $\neg B \vee \neg C$ must be inconsistent and, therefore, $B \wedge C$ must be a logical truth. However, this contradicts the assumptions given by the Nontriviality constraint. Note that none of our principles for the revision operator have been used in order to derive this result. Furthermore, if our language were more expressive and allowed for conditional subformulas of sentences, the counterexample to Nontriviality would still hold. Nontriviality *cannot* be the case once the formal background framework of Section 3 is presupposed. We suspect that this is easy to overlook because it seems natural to reason $G \subseteq K$ and $H \subseteq K$ is unproblematic to satisfy by presupposing that all the propositional sentences in G and H are members of K. But taking these subset relations for granted in the context of Ramsey-test conditionals and acceptability sets also amounts to assuming that all the *conditionals* in G and H are members of K. It is this latter and stronger assumption that is easily contradicted by further constraints on the joint consistency of the formulas in question. This analysis is more or less Levi's in [13], according to which a situation such as the one described by Nontriviality cannot be the case (see Lindström and Rabinowicz [16, Section 5.3]). Lindström and Rabinowicz describe Levi's approach in the way that only a weakened version of Monotonicity is derivable from it and that this is the reason why the triviality result is avoided. We hope it is clear from our own approach, which in many important respects overlaps with Levi's, that both the Ramsey test and its Monotonicity implication hold for acceptability sets in general, where acceptability sets contain both factual and conditionals formulas as their members; yet no contradiction follows from this. Lindström and Rabinowicz's remark that Levi does not determine acceptability conditions for nested conditionals—neither did we—leads to further interesting questions on the topic of iterated belief revision, but this latter topic should not be mixed up with the original inconsistency or triviality problems that may or may not affect the Ramsey test. After all, none of the proofs of versions of the inconsistency or triviality theorem rely on the presence of nested conditionals in acceptability sets.

If we had allowed for belief states $\mathfrak{S}$ which would exclude certain possible worlds—as described at the end of Section 4—then Nontriviality would be satisfied, but at the same time the Ramsey test would show up in this more general framework in a different form: if G is nontrivial *and* $A > \bot \notin G$, then for all factual sentences B: $B \in G * A$ if and only if $A > B \in G$. This would again preclude the derivation of a contradiction from the Nontriviality condition above together with the other principles.

7 Gärdenfors' Own Version of the Inconsistency Theorem

Gärdenfors [4; 5] states the inconsistency result in several versions, the perhaps most perspicuous one of which has the following premises (again we use our own terminology where useful).

Ramsey Test If A is consistent and G is nontrivial, then $A > B \in G$ if and only if $B \in G * A$.

Success $A \in G * A$.

Consistency If both G and A, when considered separately, are consistent, then $G * A$ is consistent.

Preservation If A is consistent with the set of factual members of G, then $G \subseteq G * A$.

Nontriviality II There are three sentences A, B, C, and there is an acceptability set G such that

(i) $A \wedge B$ is inconsistent,
(ii) $A \wedge C$ is inconsistent,
(iii) $B \wedge C$ is inconsistent,
(iv) the set of factual members of G is neither inconsistent with A nor with B nor with C.

Disjunction $G + (A \vee B) \subseteq G + A$.

Conjunction $(G + A) + B = G + (A \wedge B)$.

Gärdenfors proves that these conditions taken together are inconsistent. We have adapted the notion of relative consistency that is used in Gärdenfors' statement of

Preservation in the same way as we did in the case of Lindström and Rabinowicz' formulation; accordingly, we have restricted the Ramsey test to consistent formulas and nontrivial acceptability sets again. The principles that we have called 'Disjunction' and 'Conjunction' are stated in terms of Gärdenfors' expansion operator $+$ rather than in terms of the revision operator $*$. Since we have restricted our formal framework and our postulates to revision, we have to translate Disjunction and Conjunction into constraints on $*$. This is possible, since $G + A$ is supposed to be equal to $G * A$ if A is consistent with the set of factual members of G. If the latter is not the case, $G + A$ is intended to be the trivial acceptability set by Gärdenfors. As far as the intended reading of Disjunction is concerned, if A is inconsistent with the factual subset of G, then the Disjunction condition is trivially satisfied; if $A \vee B$ is inconsistent with the set of factual members of G, then the same holds for A and the condition is satisfied as well; accordingly for Conjunction. Thus, we may simply restate Disjunction and Conjunction in the following form.

Disjunction′ If A (and hence $A \vee B$) is consistent with the set of factual members of G, then $G * (A \vee B) \subseteq G * A$.

Conjunction′ If $A \wedge B$ (and hence A and B) is consistent with the set of factual members of G, then $(G * A) * B = G * (A \wedge B)$.

Disjunction and Conjunction are not regarded by Gärdenfors as having major importance for the proof of the inconsistency result; he does not list them in the explicit statements of the theorem but enumerates them beforehand as properties of belief expansion which "turn out to be useful in the proof" (see Gärdenfors [5, p. 158]). When he asks himself later whether perhaps "the real cause of the inconsistency is hidden among the additional assumptions," he does not consider Disjunction or Conjunction. As it turns out, the Ramsey test (with variables restricted to consistent sentences and nontrivial acceptability sets again), Success, Consistency, and Preservation are consistent with each other, and, as we have seen before, they actually follow from the principles of revision that we have introduced. The Nontriviality II condition—which figures simply as "nontriviality" in Gärdenfors [5]—is indeed absolutely harmless and expresses nontriviality. It can be added to the previously mentioned principles without giving rise to a contradiction; in fact, it is sufficient to assume that $\mathcal{L}_{\text{prop}}$ has at least two propositional variables in order to make Nontriviality II true in our framework. But *Disjunction′ and Conjunction′ do not hold* according to our framework and according to the principles of revision that we have introduced in Sections 3, 4, and 5.

Consider a belief state

$$\mathfrak{S} = \langle X_1, X_2, \ldots, X_n \rangle$$

such that $X_1 = \overline{A} \cup \overline{B} \cup \overline{C}$, where A, B, C and $G = \text{Th}(\mathfrak{S})$ instantiate the existence claim of Nontriviality II and $\overline{A} = \min_{\mathfrak{S}}([A])$, $\overline{B} = \min_{\mathfrak{S}}([B])$, $\overline{C} = \min_{\mathfrak{S}}([C])$. Then our Minimal Mutilation principle implies that $\mathfrak{S} * [A \vee B]$ is

$$\mathfrak{S} * [A \vee B] = \langle \overline{A} \cup \overline{B}, \overline{C}, X_2, \ldots, X_n \rangle,$$

while $\mathfrak{S} * [A]$ is equal to

$$\mathfrak{S} * [A] = \langle \overline{A}, \overline{B} \cup \overline{C}, X_2, \ldots, X_n \rangle.$$

By our satisfaction clauses for conditionals, $\mathfrak{S} * [A \vee B] \models (B \vee C) > B$ but $\mathfrak{S} * [A] \not\models (B \vee C) > B$. Thus, $\text{Th}(\mathfrak{S} * [A \vee B]) \not\subseteq \text{Th}(\mathfrak{S} * [A])$, which contradicts Disjunction′. For similar reasons, Conjunction′ is not satisfied. Except for the basic assumptions of our formal framework, we have only made use of Minimal Mutilation in order to contradict Disjunction′ and Conjunction′. Actually, Deconditionalization would be sufficient for deriving that $\mathfrak{S} * [A \vee B] \models (B \vee C) > B$. It is really $\mathfrak{S} * [A] \not\models (B \vee C) > B$ which follows from Minimal Mutilation; however, it would also follow from Deconditionalization plus some additional principle which ensures that after revising $\mathfrak{S}$ by $[A]$, $\overline{B}$ is still the set of minimal B-worlds, $\overline{C}$ is still the set of minimal C-worlds, and the ranks of the worlds in $\overline{B}$ are equal to the ranks of the worlds in $\overline{C}$. This weaker assumption is very plausible in itself and its plausibility is independent of the possible merits (or shortcomings) of Minimal Mutilation. For example, it is also satisfied by Spohn's revision procedure in Spohn [23].

While Disjunction′ and Conjunction′ do not follow from our postulates, it is easy to see that the propositional restrictions of Disjunction′ and Conjunction′ indeed do.

Restricted Disjunction′

> If A (and hence $A \vee B$) is consistent with the set of factual members of G, then the set of factual members of $G * (A \vee B)$ is a subset of the set of factual members of $G * A$.

Restricted Conjunction′

> If $A \wedge B$ (and hence A and B) is consistent with the set of factual members of G, then the set of factual members of $(G * A) * B$ is identical to the set of factual members of $G * (A \wedge B)$.

It is easy to mistakenly translate the restricted versions of Disjunction′ and Conjunction′ from a context where beliefs sets and belief revision are concerned with beliefs-to-be-true, to the context of belief revision for acceptability sets and Ramsey-test conditionals in which acceptability sets contain conditionals as members that have only a suppositional interpretation. While the Ramsey test is unproblematic in itself, and while it is even consistent with a principle as strong as Preservation, it cannot be combined with the Disjunction′ and Conjunction′ conditions from above. From the viewpoint of Section 2.2, the straight forward suggestion is to give up the latter rather than to abandon the former, which leaves us with a set of principles that is *almost* the one that was intended originally by Gärdenfors.

8 Segerberg's Version of the Inconsistency Theorem

Segerberg [22] derives a contradiction from three postulates on belief revision and three further, very weak, assumptions on a logical consequence operator Cn. He spells out his revision principles in terms of an addition operator $+$; however, this addition operator cannot be regarded as Gärdenfors' expansion operator $+$ from above, because Segerberg's Consistency Preservation principle, according to which if A is consistent, then $T + A$ is consistent, is not satisfied by expansion in the sense of Gärdenfors. Moreover, Segerberg's postulate of Addition Monotonicity, that is, if $T \subseteq T'$, then $T + A \subseteq T' + A$, is satisfied by Gärdenfors' expansion even without presupposing the Ramsey test, such that Segerberg's paper could not be interpreted as a note on Gärdenfors' triviality result. Fortunately, it does not do any harm to Segerberg's intended presentation of the inconsistency theorem, if we simply replace '$+$' by '$*$', since the antecedent conditions which Segerberg adds to his Expansion postulate and his postulate of Consistency Preservation ensure that addition coincides with revision in the described circumstances. Thus, Segerberg's principles can be expressed as follows.

Monotonicity If $T \subseteq T'$, then $T * A \subseteq T' * A$.
Expansion If A is consistent with T, then $T * A = Cn\{T \cup \{A\}\}$.
Consistency Preservation If A is consistent, then $T * A$ is consistent.
Consequence Operator

(R) $\Sigma \subseteq Cn\Sigma$.
(M) If $\Sigma \subseteq \Sigma'$, then $Cn\Sigma \subseteq Cn\Sigma'$.
(T) $CnCn\Sigma \subseteq Cn\Sigma$.

We have used Segerberg's original metavariables 'T' and 'Σ' instead of, say, 'G', because 'T' and 'Σ' are supposed to run over sets of *factual* formulas. This may seem surprising, because in the current context the presence of the Ramsey test acceptability sets are bound to contain conditionals as members, but presumably the underlying line of reasoning is that the Ramsey test implies Monotonicity; Monotonicity can be stated without relying on conditionals in the object language; therefore, if Monotonicity is proven to be inconsistent—given some further assumptions—then additional elaborations on conditionals can be avoided. However, it turns out that it is precisely this move which precludes the applicability of the resulting theorem to the purely suppositional understanding of the Ramsey test.

The problem is Segerberg's Monotonicity condition,

$$\text{If } T \subseteq T', \text{ then } T * A \subseteq T' * A,$$

is *not* entailed by the Ramsey test. Assume that B is a member of $T * A$: the standard argument proceeds by concluding from the Ramsey test that $A > B \in T$, which implies by $T \subseteq T'$ that $A > B \in T'$ and by the Ramsey test again that $B \in T' * A$. This proof does not go through if T and T' are purely factual formula sets. Accordingly, although

> For all factual sentences A, for all acceptability sets G, G' where G' is non-trivial:
> If $G \subseteq G'$, then the set of factual members of $G * A$ is a subset of the set of factual members of $G' * A$.

holds in our theory—this is just the Monotonicity principle from Section 4 again—the following principle is not entailed and indeed false given our framework and our postulates on revision:

> For all factual sentences A, for all acceptability sets G, G' where G' is non-trivial:
> If the set of factual members of G is a subset of the set of factual members of G', then the set of factual members of $G * A$ is a subset of the set of factual members of $G' * A$.

The latter principle is a much stronger one, which is not a consequence of the Ramsey test. Its antecedent condition, which only says something about the least-ranked layers of the two belief states in question, is extremely weak. This is in contrast with the antecedent of our own Monotonicity postulate which expresses a constraint on *all* layers of the two belief states, namely, that the second belief state is a refinement of the first one.

In Segerberg's proof the problem shows up as follows: making use of the given clauses on the consequence operator, he constructs three sets T_0, T_1, T_2 of factual formulas such that (i) $A_0 \in T_0$ and T_0 is consistent with B, (ii) $A_1 \in T_1$ and T_1 is consistent with B, (iii) $T_0, T_1 \subseteq T_2$, where (iv) A_0, A_1, B are factual formulas which are inconsistent taken together. After showing that $A_0, B \in T_0 * B$ and $A_1, B \in T_1 * B$, he concludes from Monotonicity that $A_0, A_1, B \in T_2 * B$, which would imply that $T_2 * B$ is inconsistent contrary to Consistency Preservation. While this is, of course, a perfectly sound derivation, the same argument could not be applied to acceptability sets T_0, T_1, T_2 which include conditionals and which conform to the Ramsey test and to our additional postulates: (i) would entail that $B > A_0 \in T_0$, (ii) would entail that $B > A_1 \in T_1$, (iii) would have the consequence that $B > A_0, B > A_1 \in T_2$. Whence, $B > (A_0 \wedge A_1 \wedge B)$ would have to be a member of T_2, but since A_0, A_1, B are inconsistent together by assumption, B would have to be inconsistent, which would contradict the given premises.

Segerberg leaves the interpretation of his version of the triviality result open and does not use the result in order to argue against the Ramsey test. His primary aim is to show that the proof of his formulation of Gärdenfors' theorem does not rely on strong assumptions on the logical background system. Since the premises from which he derives a contradiction are not entailed by the conjunction of the Ramsey test with other plausible principles of the belief revision of acceptability sets, the Ramsey test for conditionals remains unaffected if taken to apply to Ramsey-test conditionals.

9 Conclusions

Let us take stock. In the first two sections we took one step back from Gärdenfors' original inconsistency or triviality result on belief revision with conditionals by reformulating its assumptions in terms of an epistemic operator *Op* which might or might not be revision in the AGM sense and by replacing belief in the truth of conditionals on the right-hand side of Gärdenfors' version of Ramsey test by the mere acceptability of conditionals, where the exact meaning of 'acceptable' was yet to be determined. We discussed the three logical possibilities of avoiding inconsistency or triviality, one of which (Section 2.2) was in terms of conditionals which could be acceptable but which could not be believed to be true. For such Ramsey-test conditionals we then aimed to find out which of the standard assumptions on the revision of belief sets would have to be given up if belief sets were replaced by acceptability sets. In Section 3 we introduced the corresponding semantic representation of belief states and of the acceptability sets that they determine. We added four minimal requirements on the belief revision of acceptability sets with Ramsey-test conditionals in Section 4: the characterization of the set-theoretic type of the revision operator, the Consistency and the Inconsistency principles, and the Deconditionalization principle. These propositions are both plausible and rather weak constraints on belief revision, but they have important corollaries including the Success principle and the Ramsey test (in its purely suppositional form). Section 5 indicated one possible way of extending these minimal requirements on revision by adding the Minimal Mutilation principle. In contrast to the former, Minimal Mutilation is certainly not inevitable, although it is supported by the underlying conservativeness maxim of belief revision. We chose Minimal Mutilation mainly because Preservation (beside Cautiousness) follows from it without giving rise to any contradiction in the presence of the Ramsey test. This was in contrast with the dichotomy that very often is put forward in debates on the Ramsey test, that is, that either the Ramsey would have to be given up or Preservation. Sections 6, 7, and 8 analyzed three variants of Gärdenfors' inconsistency or triviality result. As we showed, almost all premises that have been assumed in the derivation of these results are consequences of our own principles or at least versions of them are. The only premises in 6 and 7 which do not hold according to our purely suppositional model of the revision of acceptability sets with Ramsey-test conditionals are auxiliary ones: in one case, a Nontriviality assumption that cannot be satisfied if the belief states of a cognitive agent have the formal structure of ranked models, in the other case, a postulate on disjunction and conjunction that would have to be translated unjustifiedly from a belief-to-be-true context to a context in which conditionals are merely accepted. The crucial premise which we found not to hold for the revision of acceptability sets in Section 8 is not implied by the Ramsey test and results from the exclusion of conditionals from acceptability sets. Once conditionals have been assumed to play the suppositional role (i) but not the descriptive role (ii) from the beginning of this paper, it is these auxiliary hypotheses and presuppositions which are responsible for the inconsistency or triviality results, and consequently they ought to be given up under that assumption.

Appendix: Adams' Logic

The logic whose derivability relation is sound and complete with respect to the semantics for factual and conditional formulas which we have stated in Section 3 is given by the following rules.

1. $\dfrac{}{A > A}$ (Reflexivity)

2. $\dfrac{\vdash_{\text{prop}} (A \leftrightarrow B), A > C}{B > C}$ (Left Equivalence)

3. $\dfrac{\vdash_{\text{prop}} (A \rightarrow B), C > A}{C > B}$ (Right Weakening)

4. $\dfrac{(A \wedge B) > C, A > B}{A > C}$ (Cautious Cut)

5. $\dfrac{A > B, A > C}{(A \wedge B) > C}$ (Cautious Monotonicity)

6. $\dfrac{A > C, B > C}{(A \vee B) > C}$ (Or)

The system 1–6 is often referred to as the system P (see Kraus et al. [10, pp. 189–90]). In our context, where acceptability sets contain not just conditionals but also factual formulas, we have to add (as also Adams [1] does),

7. $\dfrac{\top > A}{A}$ (Factuality I)

8. $\dfrac{A}{\top > A}$ (Factuality II).

If F is an arbitrary subset of $\mathcal{L}$, that is, an arbitrary subset of the union of $\mathcal{L}_{\text{prop}}$ and $\mathcal{L}_{\text{cond}}$, and if φ is an arbitrary member of $\mathcal{L}$—whether factual or conditional—we define that $F \vdash \varphi$ ("φ is derivable from F") if and only if there is a sequence of formulas in $\mathcal{L}$ such that the last member of the sequence is φ and each of the formulas in the sequence (i) are either members of F or (ii) are the results of applying one of the rules from above to earlier members in the sequence (in the case of Reflexivity, it is not necessary to look at earlier members). We abbreviate $\varnothing \vdash \varphi$ by $\vdash \varphi$ and we say that F is consistent if and only if $F \nvdash \bot$. If A and B are factual formulas, one can easily derive the following metarules which are demanded by Gärdenfors [5] as minimal conditions on $\vdash$.

1. If A is a truth-functional tautology, then $\vdash A$.
2. If $\vdash A \to B$ and $\vdash A$, then $\vdash B$.
3. Not $\vdash \bot$.

The corresponding semantic notion of logical consequence is given by

> If F is an arbitrary subset of $\mathcal{L}$ and φ is an arbitrary member of $\mathcal{L}$, then $F \models \varphi$ ("φ follows logically from F") if and only if for all belief states $\mathfrak{S}$: if $\mathfrak{S}$ satisfies all members of F, then $\mathfrak{S} \models \varphi$.

As Kraus et al. [10] show, for all F and φ: $F \vdash \varphi$ if and only if $F \models \varphi$. Furthermore, if we define F to be satisfiable if and only if there is a belief state that satisfies all members of F simultaneously, then one can show that F is consistent if and only if F is satisfiable.

Finally, we may define a logical closure operator Cn in the way that for all $F \subseteq \mathcal{L}$: $CnF = \{\varphi \in \mathcal{L}: F \models \varphi\}$. By the soundness and completeness result, we could have just as well defined Cn on the basis of $\vdash$. Acceptability sets G, that is, subsets of $\mathcal{L}$ for which there is a belief state $\mathfrak{S}$, such that $G = \text{Th}(\mathfrak{S})$, can be shown to be logically closed: for all $F \subseteq \mathcal{L}$, if F is an acceptability set, then $CnF = F$. The converse direction does *not* necessarily hold; but this could be either "amended" by turning from the class of ranked models to the more complex class of so-called preferential models on the semantic side or by introducing negated conditionals and the rule of "Rational Monotonicity" on the syntactic side (cf. Kraus et al. [10], Lehmann and Magidor [11]). For the same reason, it is not the case that $G * A = Cn(G \cup A)$ holds for all factual formulas A which are consistent with G, only that $G * A \supseteq Cn(G \cup A)$: in order to have the subset relation, too, one would again need to make modifications semantically or syntactically as just mentioned. In particular, if belief states are defined as preferential models rather than as linearly preordered ranked models, then a theory of belief revision can be developed that is very similar to the one we have presented in this paper such that the revision of an acceptability set G by a factual formula A that is consistent with G corresponds again to the logical closure of $G \cup A$.

Notes

1. The core of this article was written in 2005 during a research stay at Stanford. Back then its main conclusion was that one could avoid the inconsistency or triviality conclusion of Gärdenfors' [4; 5] theorem for the Ramsey test for conditionals by replacing *belief in a conditional* as expressed by the right-hand side of Gärdenfors' version of the Ramsey test by means of *conditional belief*. As Wlodek Rabinowicz rightly convinced me on his visit to Stanford, still in 2005, this more or less amounted to Isaac Levi's position in

the debate on the Ramsey test, which is why I gave up on the article for more than four years. (However, I published an independent article later titled "Beliefs in conditionals vs. conditional beliefs"; cf. [12]). In 2009, Hans Rott told me about a completely independent unpublished draft of his ("The Ramsey test for conditionals and iterated theory change") that he had been working on for quite some time and which, as it turned out, partially contained similar ideas, however, stated for the more general case of iterated belief revision. We exchanged papers, and he ended up encouraging me to believe that various details of my 2005 paper would actually be quite relevant even to current views on the Ramsey test. While the present paper is still based on my draft from 2005, I wrote a new introduction for it, and I added the completely new Section 2 on the different possible ways out of the inconsistency or triviality conclusion of Gärdenfors' theorem.

2. From Section 3 we will make the underlying language formally more precise, but in the first two sections will keep things more informal.

3. Partially, the theory of belief revision was devised to capture the effect that *learning the new evidence* A has on an agent's current belief set G. However, in the present context, the intended interpretation of $*$ is not so much concerned with learning but rather with *supposing* or *assuming* A such that $G * A$ is not the agent's *actual* new belief set that results from revising the agent's current belief set G in light of the new evidence A, but really $G * A$ is the agent's *hypothetical* or *offline* belief set that results from imposing the supposition of A to the agent's current belief set G. The difference between learning and supposing shows up very clearly whenever the antecedent and/or the consequent of a conditional describes an agent's state of mind, as in the well-known Thomason conditionals such as 'If my wife is unfaithful, I don't know it'. Learning the antecedent would allow the agent to introspect rationally that he knows that his wife is unfaithful, while merely supposing the antecedent would not; the latter option is the right one, of course, as far as the evaluation of conditionals goes, as it should be possible in principle to accept the conditional from before. With conditionals like these the agent ought to assess them by revising his beliefs in light of their antecedents while at the same time inhibiting the introspective access to his own state of mind. Fortunately, since we excluded our As and Bs to speak about an agent's internal states from the start, we will not have to worry about these special cases.

4. Lindström and Rabinowicz [16; 17] offer an excellent entry into the vast amount of literature on this topic.

5. We borrow this term from Levi [13] (see, e.g., p. 69) and, in particular, from Levi [14].

6. We should note that the AGM axioms of belief revision *qua* axioms of the epistemic operation of revision of belief have not remained undisputed themselves, even when applied to straightforwardly factual statements about the world that do express propositions and which do not contain any nonmaterial conditional sign; see, for example, Rott [19; 21]. However, often these critical remarks on AGM belief revision are orthogonal to the triviality worries about conditionals, and sometimes problems which have been observed rightly to affect the axioms of belief revision *qua* theory of revision by evidence do not apply if they are viewed as giving us a theory of a special sort of supposition; for example, while the Success Postulate 'For all G, A: $A \in G * A$' is controversial if A is meant to be new data—after all, data can be flawed and might get rejected by an agent for that reason—the postulate is perfectly innocent if taken to formalize Stalnaker's remark on suppositional reasoning in his Ramsey test quotation above, when he demands

to ‘ . . . make whatever adjustments are required to maintain consistency *(without modifying the hypothetical belief in the antecedent)*’ (our emphasis). Furthermore, by Grove’s representation result, the “logic” of AGM belief revision is simply that of linear preferential orders, which should be understood as strongly supporting the AGM axioms in view of the importance of such preferential structures in such a great number of different areas of theoretical and practical rationality (cf. Rott [20]). Indeed, the AGM postulates are really just the qualitative counterparts of the axioms for primitive conditional epistemic probability measures (Popper functions) that apply to pairs of propositions, and to the best of our knowledge no one so far has questioned the latter axioms in reaction to David Lewis’s famous triviality results on conditionals and conditional probability.

7. If this is carried out, it can be shown to lead to Ernest Adams’ well-known logic for indicative conditionals in Adams [1] which coincides with David Lewis’s logic of counterfactuals as long as one disregards nestings of conditions and the application of propositional connectives to conditionals. Indeed, that is exactly how we are going to proceed from Section 3 as far as the logic of the purely suppositional $>$ will go that we will be using then; see our Appendix for the logical details.

8. Leitgeb, “A Probabilistic Semantics for Counterfactuals” (submitted) suggests to deal with the pragmatics of subjunctive conditionals in precisely this way, however, on the quantitative-probabilistic scale.

9. This is the case since the underlying logic, Adams’, is subject to “Explosion” given contradictory factual statements.

10. An answer that is in fact prescribed by some, though not all, schemes of iterated belief revision.

References

[1] Adams, E. W., *The Logic of Conditionals: An Application of Probability to Deductive Logic*, vol. 86 of *Synthese Library*, D. Reidel Publishing Co., Dordrecht, 1975. Zbl 0324.02002. MR 0485189. 35, 49, 52

[2] Alchourrón, C. E., P. Gärdenfors, and D. Makinson, “On the logic of theory change: Partial meet contraction and revision functions,” *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–30. Zbl 0578.03011. MR 793131. 22

[3] van Benthem, J., “Dynamic logic for belief revision,” *Journal of Applied Non-Classical Logics*, vol. 17 (2007), pp. 129–55. Zbl 05640545. MR 2347619. 28

[4] Gärdenfors, P., “Belief revision and the Ramsey test for conditionals,” *Philosophical Review*, vol. 95 (1986), pp. 81–93. 22, 44, 50

[5] Gärdenfors, P., *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, A Bradford Book. The MIT Press, Cambridge, 1988. MR 956051. 22, 23, 38, 40, 41, 44, 45, 50

[6] Grahne, G., “Updates and counterfactuals,” pp. 269–76 in *Principles of Knowledge Representation and Reasoning (Cambridge, MA, 1991)*, edited by J. A. Allen, R. Fikes, and E. Sandewall, The Morgan Kaufmann Series in Representation and Reasoning, Morgan Kaufmann, San Mateo, 1991. Zbl 0765.68196. MR 1142157. 29

[7] Grove, A., "Two modellings for theory change," *Journal of Philosophical Logic*, vol. 17 (1988), pp. 157–70. Zbl 0639.03025. MR 938014. 25, 33

[8] Joyce, J. M., *The Foundations of Causal Decision Theory*, Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge University Press, Cambridge, 1999. Zbl 0941.62005. MR 1719752. 29

[9] Katsuno, H., and A. O. Mendelzon, "On the difference between updating a knowledge base and revising it," pp. 183–203 in *Belief Revision*, edited by P. Gärdenfors, vol. 29 of *Cambridge Tracts in Theoretical Computer Science*, Cambridge University Press, Cambridge, 1992. Zbl 0765.68197. MR 1173428. 29

[10] Kraus, S., D. Lehmann, and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics," *Artificial Intelligence*, vol. 44 (1990), pp. 167–207. Zbl 0782.03012. MR 1072012. 33, 35, 49, 50

[11] Lehmann, D., and M. Magidor, "What does a conditional knowledge base entail?" *Artificial Intelligence*, vol. 55 (1992), pp. 1–60. Zbl 0818.68137. MR 1167275. 33, 50

[12] Leitgeb, H., "Beliefs in conditionals vs. conditional beliefs," *Topoi*, vol. 26 (2007), pp. 115–32. Zbl 1131.03003. MR 2334968. 51

[13] Levi, I., "Iteration of conditionals and the Ramsey test," *Synthese*, vol. 76 (1988), pp. 49–81. MR 956406. 28, 44, 51

[14] Levi, I., *For the Sake of the Argument. Ramsey Test Conditionals, Inductive Inference, and Nonmonotonic Reasoning*, Cambridge University Press, Cambridge, 1996. Zbl 0935.03001. MR 1384298. 28, 51

[15] Lindström, S., "The Ramsey test and the indexicality of conditionals: A proposed resolution of Gärdenfors' paradox," pp. 208–28 in *Logic, Action and Information. Essays on Logic in Philosophy and Artificial Intelligence*, edited by A. Fuhrmann and H. Rott, Walter de Gruyter, Berlin, 1996. Zbl 0939.03019. 23

[16] Lindström, S., and W. Rabinowicz, "The Ramsey test revisited," pp. 147–91 in *Conditionals: From Philosophy to Computer Science*, edited by G. Crocco, L. Farinas del Cerro, and A. Herzig, vol. 5 of *Studies in Logic and Computation*, Oxford University Press, New York, 1995. Zbl 0824.03010. MR 1469985. 23, 38, 41, 42, 43, 44, 51

[17] Lindström, S., and W. Rabinowicz, "Conditionals and the Ramsey test," pp. 147–88 in *Belief Change*, edited by D. Gabbay and P. Smets, vol. 3 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Kluwer Academic Publishers, Dordrecht, 1998. Zbl 0937.03025. MR 1743906. 51

[18] Ramsey, F. P., *The Foundations of Mathematics and Other Logical Essays*, edited by R. B. Braithwaite, Routledge and Kegan Paul, New York, 1931. Zbl 0002.00501. MR 0120155. 22

[19] Rott, H., "Two dogmas of belief revision," *The Journal of Philosophy*, vol. 97 (2000), pp. 503–22. MR 1789860. 41, 51

[20] Rott, H., *Change, Choice and Inference. A Study of Belief Revision and Nonmonotonic Reasoning*, vol. 42 of *Oxford Logic Guides*, The Clarendon Press, New York, 2001. Zbl 1018.03004. MR 1872430. 52

[21] Rott, H., "A counterexample to six fundamental principles of belief formation," *Synthese*, vol. 139 (2004), pp. 225–40. Zbl 1106.03310. MR 2060136. 51

[22] Segerberg, K., "A note on an impossibility theorem of Gärdenfors," *Noûs*, vol. 23 (1989), pp. 351–54. MR 998093. 23, 46

[23] Spohn, W., "Ordinal conditional functions: A dynamic theory of epistemic states," pp. 105–34 in *Causation in Decision, Belief Change, and Statistics*, edited by W. L. Harper and B. Skyrms, vol. 2, Reidel, Dordrecht, 1988. 33, 40, 46

[24] Stalnaker, R. C., "A theory of conditionals," pp. 98–112 in *Studies in Logical Theory*, American Philosophical Quarterly Monograph Series 2, Blackwell, Oxford, 1968. 22

Acknowledgments

I would like to thank Hans Rott and Gerhard Schurz for their input to the paper and both the Leverhulme Trust and the Alexander von Humboldt Foundation for their general support.

Departments of Philosophy and Mathematics
University of Bristol
9 Woodland Road
Bristol BS8 1TB
UNITED KINGDOM
Hannes.Leitgeb@bristol.ac.uk