To appear in Crocco, Fariñas del Cerro, Herzig (eds.) *Conditionals and Artificial Intelligence.* *(Final version 940414.)*

The Ramsey Test Revisited*

Sten Lindström and Wlodzimierz Rabinowicz[1]

1. Background

Frank Plumpton Ramsey (1903 - 1930) — the brilliant Cambridge philosopher whose short but very intensive career started when he was still a teenager only to be interrupted by his untimely death about a decade later — made important contributions to logic and even more fundamental contributions to decision theory. As is well known, he and de Finetti — independently of each other — were the founders of the so-called subjectivist (or "personalist") approach to probability. It should be mentioned here that the first comprehensive monograph on the different aspects of Ramsey's work has only recently been published (Nils-Eric Sahlin, 1990).

The Ramsey corpus is extremely limited — only a few hundred pages.[2] But this only makes his achievement even more astonishing. On the other hand, it should make it less astonishing that the acceptability test for conditionals, which bears Ramsey's name, derives from a short footnote in one of his posthumously published papers:

> If two people are arguing 'If p will q?' *and are both in doubt as to p,* they are adding p hypothetically to their stock of knowledge and arguing on that basis about q ... If either party believes non-p for certain, the question ceases to mean anything for him except as a question about what follows from certain laws or hypotheses.[3]

His formulation of the test restricts its applicability to cases in which one is in doubt as to whether the antecedent of a conditional is true or false. The generalization of the test to all cases, independently of whether the antecedent is held in doubt, accepted, or rejected, was proposed by Stalnaker in an influential article in the late 1960s.[4] Stalnaker's formulation of the test is as follows:

> 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.

Prior to Stalnaker's article, the test seems to have been mostly ignored, although Chisholm (1946) is one exception. The official name of the test ("the Ramsey test") appears to have been coined by Harper (1976a), (1976b).

2. The Ramsey test and epistemic conditionals

The Ramsey test, in its standard modern formulation, gives the following intuitive criterion for the rational acceptance of conditionals:

> A conditional proposition "If A, then B" is (rationally) accepted in a given state of belief G just in case B should be accepted if G were revised with A as a new piece of information.

In other words, in order to evaluate a conditional statement, I ask myself what would be the case if I learned that its antecedent is true. If and only if I would *then* have reasons to accept its consequent, I *now* have reasons to accept the conditional in question.

To give a classical example, suppose that I accept the conditional

(1) If Oswald did not kill Kennedy, then someone else did.

If I now were to learn that Oswald as a matter of fact did not kill Kennedy, then I would come to believe that Kennedy was killed by someone else. That is, if I were to learn that the antecedent is true, I would accept the consequent.

It is clear, however, that the Ramsey test does not apply to all conditionals. While it might seem plausible for indicative conditionals, it does not work so well for some subjunctive ones. A person who believes, like the members of the Warren Commission, that

(2) If Oswald hadn not killed Kennedy, then no one else would have,

would most likely reject the consequent of this conditional upon learning that Oswald was innocent. Hence, the left-to-right direction of the Ramsey test does not seem to work for conditionals of type (2).

For the same reason, the Warrenite would not accept the statement

(3) If Oswald had not killed Kennedy, then someone else would have.

Hence, the right-to-left direction of the Ramsey test seems to fail for (3). For, as we noted, if he received the information that Oswald did not kill Kennedy, it would be reasonable for him, given his present beliefs, to conclude that someone else did. However, this does not give him any reason to accept the subjunctive conditional (3). The latter statement makes a counterfactual claim about the world that he does not accept. Therefore it seems that both directions of the Ramsey test fail for at least some subjunctive conditionals.

The apparent difference in the acceptance conditions between indicative conditionals like (1) and subjunctive ones like (2) and (3) suggests that there might be a semantic and not just a grammatical distinction involved here. The grammatical difference between moods (indicative versus subjunctive) seems here to be a symptom of a semantic distinction — a distinction in meaning. On the one hand, we have the *epistemic* (or doxastic) conditionals that express our dispositions to change our beliefs in the light of new information. These are the ones for which the Ramsey test appears plausible. On the other hand, there are the *ontic* conditionals that we use to make factual claims about the world. The epistemic conditionals have to do with hypothetical modifications of our *beliefs* about the world, while the ontic conditionals represent the hypotheses concerning what would be the case if the *world itself* were different — they have to do with the hypothetical modifications of the *facts* rather than with the modifications of our *beliefs* about the facts. This distinction between two kinds of conditionals parallels the well-known distinction between two kinds of probabilities: the epistemic probabilities ("credences") and the ontic or objective probabilities ("chances").

It should be added, of course, that the semantic distinction between two kinds of conditionals does not exclude the possibility that one and the same conditional sentence, in a particular context, may be used both to make a factual claim *and* to express the speaker's disposition to change his beliefs. Neither do we want to exclude the possibility that a seemingly indicative construction may be used to make an ontic claim (see conditional (4) below).

With regard to grammar, the claim that English, like Latin, contains a distinction between the subjunctive and the indicative mood is quite controversial. Dudman has vigorously argued against it in a series of influential papers.[5] According to him, the right grammatical distinction between conditionals in English should be drawn in another way, so that some "indicative" conditionals will fall into the same group as the so-called "subjunctive" ones. For example, he argues that the seemingly "indicative" sentence:

(4) If Oswald does not kill Kennedy, then no one else will,

is grammatically very different from (1). Conditional (4) belongs to the same grammatical category as the seemingly "subjunctive" sentence (2). Here, we will not try to present Dudman's way of drawing his grammatical distinction. Rather, what is relevant is his attempt to support this distinction in grammar by semantic considerations. He points out that (4), when uttered "before the event", would make the same claim as (2) makes when uttered now.

According to Dudman, conditionals like (2), (3), and (4) — he refers to them simply as "conditionals" — are arrived at "... by envisaging a developing sequence of events", "by imaginatively projecting steadily *futureward* from [the time referred to in the utterance] in a fantasy *in which* the 'if'-condition is satisfied".[6]

Conditional statements like (1), which Dudman calls "hypotheticals", are instead "arrived at by arguing from proposition to proposition ...". "Someone shot Kennedy, Oswald did not, therefore someone else did."[7]

It seems clear that Dudman's semantic interpretation of his grammatical distinction between "hypotheticals" and "conditionals" has some affinity with our semantic distinction between epistemic and ontic conditionals.[8]

3. Belief revision

As we have seen, the Ramsey test is inappropriate for ontic conditionals. But even if we restrict our attention to epistemic conditionals, as we shall do in this paper, the test turns out to be problematic. Gärdenfors has shown that the Ramsey test — or at least a formal version of it — is incompatible with certain plausible conditions on *belief revision*. Gärdenfors presents this paradoxical result in (1988). His axioms for belief revision that underlie the argument are also presented in an already classical paper which he wrote with Alchourrón, and Makinson (Alchourrón, Gärdenfors, Makinson, 1985). Alchourrón and Makinson originally studied changes of norm systems (such as changes in law when a new rule is being introduced), but they soon discovered that norm revision and belief revision are processes that exhibit striking formal similarities.

Alchourrón, Gärdenfors and Makinson (AGM) identify a state of belief with the *set* of all beliefs accepted in that state, and they represent beliefs linguistically, by *sentences.* On their approach, therefore, a state of belief G is nothing but a belief set, which in turn is nothing but a certain set of sentences.

In addition, they make strong idealizing assumptions about the agents whose beliefs are being considered. In particular, they assume that the belief sets are closed under logical consequence.[9]

It is also assumed that belief revision is a *function* that transforms states of belief into new states of belief, given new information. If G is the original state and the sentence A represents the new information, then the new state of belief can be denoted as G∗A. We refer to this new state as G *revised by* A. Note that A — the new information — may or may not be logically compatible with G. In the former case G∗A may be seen as an expansion of G with A, while in the latter case G∗A must involve a genuine revision of G. Some of the old beliefs must be given up in order to make room for the new information.

It should be mentioned that some of the strong idealization assumptions that characterize the AGM approach may well be questioned. In particular, in other papers we have suggested that belief revision should be seen as a relation rather than as a function (Lindström and Rabinowicz 1989, 1992, 1990). At least when A is logically incompatible with G, the new state of belief might not be uniquely determined by the old state and the new information; there might be several reasonable ways of revising G with A. In order to incorporate A, the agent has to give up some of his original beliefs. But he can solve this problem in different ways.

We show in (1992) that the extra resources available in the relational approach can be utilized to avoid Gärdenfors' paradox. However, we shall not pursue this line of reasoning here. Instead, we shall consider the paradox in its original functional setting only.

Let us introduce the symbol > for the epistemic if-then connective. "A > B" should be read as "If A, then B". Within the AGM approach, the Ramsey test takes the following form:

(RT) For every belief set G,

$A > B \in G$ if, and only if, $B \in G{*}A$.[10]

Gärdenfors puts forward a number of postulates that belief revision should satisfy.

(Success) $A \in G{*}A$.

(Consistency) If both G and A, when considered separately, are logically consistent, then G∗A is logically consistent.

(Preservation) If A is logically consistent with G, then $G \subseteq G{*}A$.

According to Preservation, you keep your original beliefs after the revision, unless you *have* to give up some of them in order to avoid inconsistency with the new information.

There are further, more or less reasonable, postulates that one might impose on belief revision, but these three suffice to show that the Ramsey test cannot be upheld.

4. Gärdenfors' paradox

The Ramsey test implies that belief revision has to be a "monotonously increasing" (or at least non-decreasing) function:

(Monotonicity) If a belief set G is included in another belief set K, then G∗A is included in K∗A.

Proof: Suppose that

(i) $G \subseteq K$.

In order to prove that the same relation of inclusion obtains between G∗A and K∗A, we have to show that every belief B that belongs to G∗A must also belong to K∗A. Suppose, therefore, that

(ii) $B \in G{*}A$.

Applying (RT) to (ii), we obtain:

(iii) $A > B \in G$.

Therefore, given (i),

(iv) $A > B \in K$.

But then, applying (RT) once again, this time to (iv), we obtain the conclusion that we have been after:

(v) $B \in K{*}A$.

Let us now move to the second part of the argument. We have seen that the Ramsey test implies Monotonicity. Now we have to show that Monotonicity is incompatible with the assumed postulates on belief revision.

The argument assumes that there could exist two sentences B and C and three consistent belief sets G, H and K such that:

(1) $B \in G$ and $G \cup \{\neg C\}$ is consistent;

(2) $C \in H$ and $H \cup \{\neg B\}$ is consistent;

(3) $G \subseteq K$ and $H \subseteq K$.

We call this assumption *Non-Triviality*.

To illustrate the situation, let B be the sentence "It is raining in Uppsala" and C the sentence "It is snowing in Lund". Suppose that our agent starts from a belief state in which he has no opinion about the weather in either Uppsala or Lund. Let G, H and K be the states of belief that he would reach upon learning B, C and $B \wedge C$, respectively. It is reasonable to assume that such G, H and K are going to satisfy conditions (1), (2), (3).[11]

Let A be the sentence

$$\neg B \vee \neg C.$$

That is, A amounts to the claim that it does not rain in Uppsala or it does not snow in Lund.

It follows from (1) and (2) that each of G and H is logically compatible with A. Since B and C belong to G and H respectively, Preservation implies that

(i) $B \in G*A$,

(ii) $C \in H*A$.

By Success, we also know that

(iii) $A \in K*A$.

Since A is incompatible with B and C taken together, (iii) by Consistency implies that

(iv) $B \notin K*A$ or $C \notin K*A$.

As we remember, Consistency ensures that the revision of a consistent state of belief, such as K, with an independently consistent piece of information must itself be consistent.

However, since $G, H \subseteq K$ (condition 3), Monotonicity implies that $G*A, H*A \subseteq K*A$. Hence, by (i) and (ii), we are led to the denial of (iv):

(v) $B \in K*A$ and $C \in K*A$.

We have thus reached a contradiction. The postulates of Success, Consistency and Preservation, taken together, conflict with Monotonicity, and thereby with the Ramsey test.

5. In search of a solution

What should we do about this impossibility result? Something has to give. It seems that there are only two alternatives: either to reject the Ramsey test or to give up some of the postulates on belief revision. (Later on we shall also consider the possibility that the dilemma might be misleading — that it might be due to hidden assumptions in the AGM framework.)

5.1 Questioning preservation

Let us first consider the second alternative. Which postulates on belief revision could be given up? Success and Consistency seem to be relatively noncontroversial conditions. Therefore, the natural choice would be to question Preservation. Is it so obvious that a rational belief revision has to be preservative? Well, the answer partly depends on what one means by *beliefs*. According to one common view, B is one of my beliefs if I assign to B a *sufficiently high probability*. Given this interpretation, there is no reason to assume that belief revision will always be preservative. Suppose that my probability for B is initially very high, and that I subsequently receive new information that is logically consistent with B and the rest of my original beliefs (where "belief" is taken as now understood). This information, however, may well make B less probable, so that my new probability for B will no longer be "sufficiently high". In consequence, B will not belong to my new beliefs and Preservation will be violated. To illustrate, think of a district attorney who believes that the suspect is guilty, but then stumbles upon unexpected new evidence which, while insufficient to prove the suspect's innocence, still makes his guilt less probable. The new evidence is logically consistent with the district attorney's original beliefs but it still makes him give up his belief in the suspect's guilt.

Interpreting belief as sufficiently high probability is, however, inconsistent with Gärdenfors' assumption of logical closure for belief sets. The conjunction of two sufficiently probable beliefs may not itself be sufficiently probable. We need an interpretation of belief that makes belief sets closed under conjunction.

There is another interpretation of belief that may accomplish the same goal. A person's beliefs may be identified with those propositions that are *well-established* for that person, in terms of his available evidence. It may well be claimed that if A is well-established and B is well-established, then the same applies to their conjunction. However, this interpretation — just as the "sufficiently high probability"-approach — would make Preservation unreasonable. Just think of our district attorney for whom, on the evidence that he has gathered, the suspect S's guilt is well-established, but who subsequently receives new evidence which points in another direction. The new evidence may well be logically compatible with the things that the district attorney has originally considered to be well-established, including the hypothesis that S is the guilty person. However, given this new evidence, the hypothesis may cease to be well-established.

The same argument applies to what might be called a pragmatic interpretation of belief according to which the agent's beliefs in a situation are the propositions that in that situation he takes for granted and is prepared to act on (where 'acting' may also involve such epistemic activities as reasoning or planning). Let us refer to such opinions as *assumptions* or *presuppositions*. Under ideal circumstances, it may be claimed that the set of such assumptions is logically closed. In particular, for a reasonable agent, if he is prepared to act on A and on B, separately, then he is also prepared to act on A ∧ B. Again, however, Preservation is violated by this interpretation of belief. Our attorney is originally prepared to act (go to court, for example) on the hypothesis that S is guilty. But he would no longer be prepared to act on this hypothesis if he received an additional piece of evidence E that would point in another direction. However, he is not originally prepared to act on the hypothesis ¬E. That is, the assumption that S is guilty would be given up if the agent learned E, even though E is consistent with his original assumptions (cf. Rabinowicz (in press) for further discussion of the relationship between Preservation and the pragmatic interpretation of belief).

In their discussion of non-monotonic logic, Gärdenfors and Makinson (1994) introduce a notion of *expectation* that seems closely related to our concept of assumption. They start out with the notion that one proposition is *more expected* (to be true) than another. An expectation is defined as a proposition that is more expected than its negation. Gärdenfors and Makinson think that non-monotonic inference can be understood in terms of a hypothetical revision of expectations: A non-monotonically implies B if we would expect B if we came to expect A. However, they assume that revision of expectations does satisfy Preservation. This suggests that their notion of expectation is different from our assumptions. Probably, when confronted with our example of the district attorney, Gärdenfors and Makinson would say that the new information E received by the attorney is incompatible with his original expectations. Thus he originally expects ¬E although he is not prepared to act on this supposition. Possibly, then, the relationship between assumptions and expectations can be described as follows: assumptions are those among our expectations that in a given context are sufficiently well-expected in order to be taken for granted for purposes of reasoning and action.

In his (1988), Gärdenfors tends to interpret beliefs in still another way. To believe that B is to be *certain* that B is the case. It is to assign to B the highest possible degree of probability: *probability 1*. Gärdenfors may now defend Preservation by deriving it from the so-called conditionalization model for probability revision. The model in question establishes the following connection between the prior probabilities and the probabilities based on new information:

> The probability $P_A(B)$ to be assigned to B upon learning A (as the sole piece of new information) equals the original conditional probability for B given A, P(B/A). That is, by the definition of conditional probability, it equals P(A ∧ B)/P(A). Provided, of course, that the prior probability P(A) > 0. (Otherwise, the ratio would not be well-defined.)

Now, the conditionalization model implies the following thesis for the special case in which B's prior probability is equal to 1:

> If the prior probability for B equals 1 (and the prior probability for A is positive), the probability to be assigned to B upon learning A still equals 1.

In other words, our certainties should be kept intact when we receive new information that is consistent with the propositions that we have been certain of. Revision of certainties satisfies Preservation.

A very influential argument for the conditionalization model due to David Lewis (cf. Teller 1973) consists in a proof that an agent who tends to violate the model is vulnerable to a *diachronic Dutch book*: his probability assignments make him accept certain bets, over a period of time, which are so constructed that, if he takes them all, he is bound to lose whatever the outcome! (Some of the bets would be offered to him by a clever bookmaker before the arrival of the new information, while others would be kept in abeyance to be offered if and when the new information is due to arrive, if it is going to arrive at all.)

Diachronic Dutch books are natural extensions of the concept of a *synchronic Dutch book*: a system of bets simultaneously offered to an agent such that, if he takes them all, he is bound to lose whatever happens.

As an aside, it might be mentioned that the following result is due to Ramsey: at a given occasion, an agent is vulnerable to a synchronic Dutch book just in case his probability assignments on that occasion violate the axioms of the probability calculus. David Lewis has extended Ramsey's result to the diachronic case, thereby making it possible to argue that the conditionalization model has a comparable status to the standard axioms of probability.

We are not sure whether this argument for Preservation is convincing. At least prima facie it might seem that in some situations new evidence may well undermine one's original certainties without being logically inconsistent with them. Thus, for example, if one is initially certain that

> Most F's are G's

(say, that most swans are white), while lacking a definite opinion about the size of the F-population (how many swans are there, actually?), one's certainty may well be undermined by repeated observations of F's that are not G's. (Here is a black swan, and here is another one, ...) Note that such repeated observations are logically consistent with one's initial certainties, but they still may make one wonder. Clearly, starting to wonder violates the conditionalization model. Agents of this type are vulnerable to diachronic Dutch books, but it may be questioned whether such a diachronic vulnerability is a sure sign of irrationality. Is it so unreasonable to be influenced by new information in such a way that one is led to respond to it in a way that violates the conditionalization model? If one expects oneself to respond to new information in this way, then one had better avoid clever bookmakers, but there may be few clever bookmakers around anyway! (For a sustained criticism of diachronic Dutch book arguments, see Christensen, 1991.)

If these objections are well taken, Preservation would have to be rejected as a general rationality requirement on belief revision — even if we interpret belief as certainty.

On the other hand, the conditionalization model is an extremely elegant and powerful tool for describing belief dynamics. It constitutes a central component of the very lively Bayesian tradition in decision theory. It might therefore be of interest to suppose, at least for the arguments sake, that Preservation (Conditionalization) is defensible and consider the other horn of the dilemma: giving up the Ramsey test. How much of the idea behind the test can we then preserve?

5.2. Weakening the Ramsey test

There are many ways in which the Ramsey test could be modified so as to avoid Gärdenfors' paradox. It is ironical that Ramsey himself only proposed a weak form of the Ramsey test, restricted to the conditionals whose antecedents are epistemically open. That version is certainly compatible with Preservation. But it does not give us an acceptability condition for the interesting case when the antecedent of the conditional is incompatible with the agent's beliefs.

Here, we shall sketch a solution that we have proposed in Lindström and Rabinowicz (1992). There we suggest that the Ramsey test should be replaced by what we call the *strict Ramsey test:*

(Strict RT) For every belief set G,
$A > B \in G$ if and only if for every extension H of G, $B \in H*A$.

Here, by an extension of G, we mean any possible belief set H that includes G.

According to Strict RT, in order for an epistemic conditional to be accepted with respect to a given belief set G, it is not enough that the conditional's consequent should be accepted if G were revised with the antecedent. It is also important to consider what would happen if G were first expanded in different ways. Should the consequent still be accepted if such an extension were revised with the antecedent? Thus Strict RT replaces the condition

(1) $B \in G*A$

by the logically stronger condition

(2) for every extension H of G, $B \in G*A$.

At the same time, it preserves Ramsey's principal idea: the acceptance criterion for (epistemic) conditionals is still formulated in terms of hypothetical changes of belief.

It is easy to show that the original Ramsey test implies the strict version. We assume RT and try to prove Strict RT. The right-to-left direction of the latter condition immediately follows from the former in view of the fact that G is one of its own extensions. To prove the left-to-right direction, note that if A > B ∈ G, then for any extension H of G, A > B ∈ H. But then, applying RT from left to right, it follows that B ∈ H∗A.

Strict RT in turn implies the following condition:

(Normality) For every *opinionated* belief set G,
A > B ∈ G if, and only if, B ∈ G∗A.

An opinionated belief set is a belief set which is complete, or maximally consistent, in the sense of containing for every sentence A, exactly one of the sentences A and ¬A. Normality follows from Strict RT in view of the fact that an opinionated belief set has no proper consistent extensions.

Normality is just the original Ramsey test restricted to opinionated belief sets. This condition is equivalent to what we call the *weak Ramsey test:*[12]

(Weak RT) For every belief set G,
A > B ∈ G if, and only if, for every opinionated extension H of G, B ∈ H∗A.

Thus we have: RT ⇒ Strict RT ⇒ Weak RT. In fact, Strict RT is equivalent to Weak RT (or Normality) together with the following condition on the revision operation:

(Connection Downwards)
For every consistent belief set G and every A,
$\bigcap\{H*A$: H is an opinionated extension of G$\} \subseteq G*A$.

This condition implies that the right-hand sides of Strict RT and Weak RT are equivalent.

It is natural to divide the Ramsey test into its 'only if' part

(RT⇒) A > B ∈ G, only if B ∈ G∗A

and its 'if' part

(RT⇐) A > B ∈ G, if B ∈ G∗A.

As we can easily prove, the strict Ramsey test is equivalent to the weak Ramsey test (or Normality) together with *(RT⇒)*. In the presence of Normality, *(RT⇒)* is equivalent to Connection Downwards. We can also prove that the Ramsey test is equivalent to the strict Ramsey test together with *(RT⇐)*. Furthermore, in the presence of Strict RT, *(RT⇐)* is equivalent to Monotonicity.

Strict RT and Monotonicity are two logically independent conditions, of which only the latter is involved in the derivation of Gärdenfors' paradox. One way of avoiding the paradox is to replace the Ramsey test by the strict Ramsey test. That is, the suggestion is to keep the Ramsey test for opinionated belief sets (Normality) together with the 'only if' part of the full Ramsey test, while giving up the 'if' part of the test.

It is compatible with this proposal to think of epistemic conditionals as having not only acceptance conditions but also truth conditions: they may be thought of as bearers of truth values. Actually, the truth condition of a conditional may be seen as its acceptance condition with respect to an appropriate set of beliefs. Using the language of possible worlds, we may say that to each possible world w corresponds the set of all beliefs that would be true at w. Let us refer to that set as G_w. Note that G_w is a logically consistent set of beliefs (otherwise w would not be a possible world) and that, for each potential belief A, either A or its negation belongs to G_w. Thus G_w is an opinionated belief set; it cannot be consistently expanded. Obviously, a conditional A > B is true at a possible world w if and only if it belongs to G_w. By Normality, A > B belongs to G_w if, and only if, B belongs to G_w*A. (Recall that Strict RT implies Normality.) Thus we reach the following *truth condition for epistemic conditionals:*

> A > B is true at a possible world w if, and only if, B belongs to the revision of the associated belief set G_w with A.

Using the idea, going back to Carnap, of identifying the proposition expressed by a sentence with the set of possible worlds in which the sentence is true, we can now associate with each conditional sentence A > B the proposition that this sentence expresses: the set of worlds w such that $B \in G_w*A$.

Why does the strict version of the Ramsey test allow us to avoid Gärdenfors' paradox? Well, the answer is simple. It is no longer possible to derive Monotonicity. Suppose we assume (i) that $G \subseteq K$ and (ii) that $B \in G*A$. As we remember, to prove Monotonicity, we need, as a first step, to derive from (ii) the intermediate conclusion (iii) that $A > B \in G$. Then we will be able, using Strict RT, to derive from (i) and (iii) the final conclusion (iv) that $B \in K*A$. However, given Strict RT, (iii) no longer follows from (ii). In order to derive (iii), we need a stronger assumption. It is not enough that B belongs to G∗A. B must also belong to H∗A, for every extension H of G. But K itself is one of the extensions of G (according to the assumption (i)). Therefore we would have to assume (iv) — the final conclusion that we are after — in order to derive the intermediate conclusion (iii). Clearly, then, the proof does not go through.

Thus the paradox is avoided. But is Strict RT a reasonable replacement for RT? One might object against Strict RT and say that this principle is *too* demanding. If our acceptance of a conditional in a belief state would be dependent on *all* the extensions of that state, then perhaps we would very seldom be in a position to accept a conditional sentence. One could also point out that, in order to determine whether a conditional should be accepted, there is no reason to consider extensions that are too outlandish to be taken seriously.

However, this criticism is not really well taken, at least from the point of view of someone who accepts the 'only if'-part of the Ramsey test (RT⇒). (Remember that this condition is accepted both by the adherents of RT and by those who want to replace it by Strict RT.) For consider what it would involve to construct a counterexample to Strict RT. We would need a belief set G such that A > B is accepted in G, together with an extension K of G such that $B \notin$

K∗A. But a moment's reflection suffices to realize that this is impossible. By (RT⇒), $B \notin K*A$ entails that $A > B \notin K$. But then K cannot be an extension of G, since by hypothesis $A > B \in G$.

There is an air of paradox in this defence of Strict RT. One has the feeling that counterexamples might still be produced. Think of someone who is confronted with a match that to all appearances is perfectly normal. Given his present belief state G, he would believe that the match would light upon learning that it is going to be struck. Suppose, however, that instead he learns a new fact: there is a man hidden in the vicinity with a fire-extinguisher trying to prevent the match from catching fire. The addition of such an "outlandish defeater" moves him to a new belief state K in which he would no longer expect the match to light, if he *now* learned that it is going to be struck. It seems that the defeater is consistent with G, because of its outlandish character — the agent has never considered such a possibility, and so he has never had a need to exclude it. Hence, it seems that K is an extension of G. At the same time it seems reasonable to claim that in G the agent accepted the conditional: "If the match is going to be struck, then it will light". So we do seem to have a counterexample to Strict RT.

However, if the conditional in question was accepted in G, then K cannot be an extension of G because this conditional is no longer accepted in K (in virtue of RT⇒). Note that our notion of an extension is very strong: in order for K to be an extension of G, K must contain all the sentences that occur in G, including the conditional ones. We could, of course, also consider a weaker notion of extension. Say that $\mathcal{L}_0$ is the set of all "factual" sentences of the object language, that is, the sentences that do not contain any occurrence of the conditional connective > or any other modal constructions (to be more precise, $\mathcal{L}_0$ consists of all the sentences that can be constructed from the atomic sentences of our object language by means of the standard logical operations). Then

> H may be said to be a *factual extension* of G (in symbols, $G \sqsubseteq H$) if every $\mathcal{L}_0$-sentence A in G is also a member of H.

Clearly, all extensions of G are also factual extensions, but not vice versa. In particular, the belief state K in our fire-extinguisher example, while not an extension of G, is presumably one of its factual extensions.

Now, perhaps those who think Strict RT is unacceptably demanding confuse it with a parallel test formulated in terms of factual extensions:

(Super-Strict RT) For every belief set G,
$A > B \in G$ if, and only if, for every *factual* extension H of G, $B \in H*A$.

Clearly, Super-Strict RT, which is much more demanding than Strict RT, is vulnerable to the kind of counterexamples described above. But Strict RT still seems to be a plausible Ramsey-type test.

5.3. The non-propositional interpretation of conditionals

Up to now, we have been considering "conservative" solutions of Gärdenfors' paradox — solutions that do not question its underlying conceptual framework. But surely this framework could be criticized. Thus, in Levi's view (Levi, 1988), the paradox crucially depends on Gärdenfors' assumption that conditional sentences are members of belief sets. In other words, the assumption is that conditionals express truth-value bearing propositions and may therefore be statements of belief. Instead, Levi claims that conditionals, rather than stating beliefs, are appraisals of beliefs with respect to their epistemic possibility. Rather than being true or false claims about reality, they are *epistemic evaluations*. According to Levi,

> a corpus of knowledge [a belief set] is a resource for deliberation and is to be analyzed in terms of its function as such a resource ... And the chief function [of] a corpus of knowledge is to serve as a standard of epistemic possibility. If [B] is in the corpus, its negation is not a serious possibility. If [B] is not in the corpus, its negation is a serious possibility. (ibid., p. 55)

Now, expressions such as "It may well be that ..." are used to express our evaluations of serious possibility, based on our current belief set. The role of conditionals is different but related:

> The key idea [...] is this: a conditional of the (regimented) type [A > B] is a judgement concerning the serious possibility of [B] relative to a *transformation* of the current corpus or belief set [...] and not relative to the current corpus itself. (ibid., p. 61)

To be more precise, we evaluate the consequent of a conditional basing ourselves not on our current belief set but on the potential *revision* of it with the antecedent. To accept a conditional A > B is to appraise ¬B as not seriously possible if judged on the basis of the revision of the current belief set with A.

Clearly, such an interpretation of conditionals immediately leads to something like the Ramsey test. But Gärdenfors' formulation of this test — as a principle specifying the condition under which a conditional is a *member* of a belief set — is no longer viable.

On Levi's view, if G is my current belief set, then the conditionals that are accepted with respect to G do not themselves belong to G but rather to the associated corpus RL(G) of epistemic appraisals. Belief sets are sets of sentences of the 'factual' language $\mathcal{L}_0$. RL(G) may be defined as the smallest set of sentences of the object-language $\mathcal{L}$ that satisfies the conditions:

(i) $G \subseteq RL(G)$

(ii) if $B \in G{*}A$, then $A > B \in RL(G)$

(iii) if $B \notin G{*}A$, then $\neg(A > B) \in RL(G)$

(iv) RL(G) is closed under truth-functional consequence.

(As a matter of fact, Levi assumes that RL(G) also contains 'unconditional' epistemic appraisals such as "It may be the case that B". Also, Levi allows that RL(G) may be closed under first-order logical consequence and not just under truth-functional entailment. But let us ignore these complications here.)

Thus, RL(G) contains all the beliefs of the agent together with the sentences that express the agent's conditional epistemic evaluations of his beliefs. Among the sentences of RL(G), only the sentences belonging to G, the belief set proper, are true or false.

Define language $\mathcal{L}_1$ as the smallest set of sentences that is closed under standard truth-functional operations and that contains all the sentences of the factual language $\mathcal{L}_0$ together with all the conditional sentences of the form A > B, where A, B are sentences of $\mathcal{L}_0$. Note that the sentences in RL(G) all belong to $\mathcal{L}_1$.

A sentence A of $\mathcal{L}_1$ is said to be *accepted* in G if and only if A belongs to RL(G). For a logically consistent G, the definition of RL(G) now implies the following version of the Ramsey test:

(Levi's RT) For any sentences A, B of $\mathcal{L}_0$ and for any belief set $G \subseteq \mathcal{L}_0$, A > B is accepted in G if and only if $B \in G*A$.

Note also that we obtain an analogous test for the negated conditionals:

(Levi's negative RT) For any sentences A, B of $\mathcal{L}_0$ and for any belief set $G \subseteq \mathcal{L}_0$, ¬(A > B) is accepted in G if and only if $B \notin G*A$.

The difference between RT and Levi's RT may seem negligible but it has far-reaching consequences. From Levi's RT it is impossible to derive *Monotonicity*. It is impossible to prove that

for all consistent belief sets $G, K \subseteq \mathcal{L}_0$ and all A in $\mathcal{L}_0$, if $G \subseteq K$, then $G*A \subseteq K*A$.

Thus, Gärdenfors' paradox is dissolved.

Monotonicity can no longer be derived because, given Levi's RT, the assumptions from which the standard derivation starts,

(i) $G \subseteq K$

and

(ii) $B \in G*A$,

only imply that

(iii) $A > B \in RL(G)$.

To prove Monotonicity, we would now have to derive

(iv) $A > B \in RL(K)$

from (i) and (iii). Then another application of Levi's RT would give us the desired conclusion: $B \in K*A$. But (iv) does not follow. To derive (iv) from (iii) we would need to replace (i) by a considerably stronger assumption:

(i') $RL(G) \subseteq RL(K)$.

In other words, we avoid Gärdenfors' paradox, because Levi's RT only allows us to prove a *weak* version of Monotonicity:

(Levi's Monotonicity) For all consistent belief sets G, K $\subseteq \mathcal{L}_0$ and all A in $\mathcal{L}_0$,
if RL(G) $\subseteq$ RL(K), then G∗A $\subseteq$ K∗A.

At this point someone might object and point out that such a weak form of monotonicity may be sufficient to generate paradoxical consequences. For Levi is prepared to accept Gärdenfors' postulates on revision of belief sets (interpreted by Levi as sets of $\mathcal{L}_0$-sentences), provided that we restrict the postulates to revision with $\mathcal{L}_0$-sentences. Thus, he is prepared to accept that the revision of any belief set G with an $\mathcal{L}_0$-sentence A is successful, that it is consistent if G and A are separately consistent, and that it is preservative if A is consistent with G. But then, why can we not use these postulates together with Levi's Monotonicity to construct a new version of the paradox? The only thing that we have to do is to strengthen the initial assumption of the argument. As we remember, the derivation of the original paradox started from the assumption that there could exist sentences B and C and consistent belief sets G, H, and K such that (1) G contains B but not ¬C, (2) H contains C but not ¬B, *and* (3) G, H $\subseteq$ K. To derive the paradox, we considered the revisions of the three belief sets in question with A = ¬(B ∧ C). Since B and C were 'factual' $\mathcal{L}_0$-sentences (B = it is raining in Uppsala, C = it is snowing in Lund), the same must apply to A. Therefore, the revision with A must obey the $\mathcal{L}_0$-versions of Gärdenfors' revision postulates. Thus, in order to reintroduce the paradox within Levi's framework, we only need to replace (3) with a stronger assumption:

(3') RL(G), RL(H) $\subseteq$ RL(K).

Then Levi's Monotonicity can be put to use in the derivation of the paradox.

However, Levi would have an easy answer to this objection. One person's modus ponens is another person's modus tollens. The situation envisaged is impossible! By Preservation, (1) and (2) imply that B is still in G∗A while C is in H∗A. Therefore, by Levi's RT, A > B must be in RL(G), while A > C must be in RL(H). On the other hand, by Success and Consistency, A is in K∗A and K∗A is consistent. Therefore, one of the sentences B and C is not in K∗A, so that — by Levi's RT — one of the conditionals A > B and A > C is not in RL(K). Consequently, RL(K) cannot include both RL(G) and RL(H). Assumption (3'), unlike (3), cannot obtain. Even though G and H may both be included in K, some of the conditionals accepted in G or H will have to disappear as we move to K. The impression that (3') could obtain arises only because we do not clearly distinguish between a belief set and the set of sentences that are accepted in this belief set.

Actually, given Levi's *negative* RT (which follows from his definition of RL), an extremely weak postulate on belief revision is sufficient to prove an even stronger result. The situation envisaged is impossible, because there cannot exist consistent belief sets G and K such that K is a *proper* extension of G but RL(G) is included in RL(K).

(Collapse) For all consistent belief sets G and K, if RL(G) $\subseteq$ RL(K), then G = K.

To prove Collapse we only need to postulate that revising by a tautology does not change our beliefs. That is, letting T stand for any tautology in $\mathcal{L}_0$, we only need to assume the following condition:

(T) For all consistent belief sets G, G∗T = G.

We can now prove Collapse using Levi's negative RT. First, we notice that RL(G) ⊆ RL(K), immediately yields G ⊆ K. Suppose, for *reductio*, that the converse inclusion does not hold. Then, for some B ∈ K, B ∉ G. By T, B ∉ G∗T. Therefore, by Levi's negative RT, ¬(T > B) ∈ RL(G), so that this negated conditional must also belong to RL(K). Applying Levi's negative RT once again, we obtain: B ∉ K, which conflicts with the *reductio* hypothesis. Q.E.D. (That the negative RT together with T lead to a collapse was originally proved by Gärdenfors *et al.*, 1990.)

To conclude, then, Levi's diagnosis is clear. The apparent paradox rests on a confusion between beliefs and epistemic evaluations. It dissolves when we interpret conditionals in a non-propositional way and thereby remove them from belief sets.

As it stands, this solution does not give all the answers we need. In particular, Levi's RT does not determine acceptance conditions for *nested* conditionals, in which some occurrences of the conditional connective appear within the scope of its other occurrences. (Note that nested conditionals are not among the sentences of $\mathcal{L}_1$. Thus, they never appear in the RL-set associated with a given belief set.) Among the nested conditionals, the *iterated* ones are of special interest. Here, it is enough to consider the simplest cases of iteration:

(i) A > (B > C) and (ii) (A > B) > C.

We take A, B, and C to be sentences of $\mathcal{L}_0$. The following conditionals seem to instantiate (i) and (ii), respectively: "If the train leaves on time, then I will miss it if I don't rush"; "If the watch will be damaged if I use it when diving, then I have been cheated".

If conditionals are interpreted as conditional appraisals of beliefs, as Levi does, then it might seem that expressions instantiating (i), and perhaps even (ii), must be judged to be meaningless. For example, in (i), what seems to be conditionally appraised is not a possible belief but an appraisal of belief — another conditional! The trouble is, however, that iterated conditional constructions do seem to be used by us, so that — as Levi admits — they cannot be totally devoid of meaning. At least, we must be able to determine their acceptance conditions.

It is easy to do it, as far as conditionals of type (i) are concerned. A > (B > C) should be accepted in a belief set G if and only if C ∈ (G∗A)∗B. The acceptance condition for such an iterated conditional is framed in terms of iterated revision.

Conditionals of type (ii) cannot be dealt with in this manner. Nor can we say that (A > B) > C is to be accepted in G iff C ∈ G∗(A > B). According to Levi, belief sets can only be revised with "factual" sentences — members of $\mathcal{L}_0$. He suggests, therefore, another solution (ibid., p. 76f). Generally speaking, (A > B) > C should be accepted in G iff C ∈ H, where the belief set H is an appropriate transformation of G — a transformation in which the antecedent of the iterated conditional would be accepted. Intuitively, H is a belief set that satisfies the following condition: A > B is accepted in H (i.e., by Levi's RT, $B \in H*A$) and H otherwise differs from G as little as possible.

What sort of transformation are we here talking about? Levi is not sure whether there is a unique answer to this question. In some cases, however, but perhaps not always, the transformation in question will consist in revision: H will be the result of revising G with some appropriate "factual" sentence D, i.e., $H = G*D$. Intuitively, D constitutes what might be called the *(potential) ground* for A > B relative to G. If you want to identify D, you should ask yourself what "factual" belief would be necessary and sufficient for the acceptance of A > B. Levi suggests that the belief in question, in some cases at least, might consist in an ascription of a dispositional property to an object or a system of objects. Thus, in the "watch"-example above, D might be the sentence "The watch is disposed to be damaged on being used when diving". Provided, at least, that such dispositional sentences are "factual" — as Levi takes them to be. It won't do if they themselves are to be analyzed in terms of conditionals.

Is it always possible to find a dispositional sentence that constitutes the ground for a given conditional? Levi is unclear on this point. However, it seems that finding such "dispositional" grounds for some conditionals might not be easy, to say the least. Thus, for example, what dispositional sentence could ground the Oswald-Kennedy conditional "If Oswald did not kill Kennedy, then someone else did"? Surely, sentences such as "Kennedy was disposed to be killed by someone else than Oswald on not being killed by Oswald", or "Someone else than Oswald was disposed to kill Kennedy on Kennedy not being killed by Oswald" are too grotesque to be even considered as possible candidates.

It seems, therefore, that Levi's approach to iterated conditionals leaves us with many unanswered questions. What other sentences are there, besides the dispositional ones, that could constitute grounds for conditionals? And what are we to do if a given conditional lacks a grounding sentence at all? How are we to find the appropriate transformation H of G in which the antecedent conditional is accepted, if we cannot assume that H must always be a revision of G with some factual sentence?[13] We are being left in the dark.

5.4. The indexical interpretation of conditionals

Isaac Levi pointed out a tacit assumption behind Gärdenfors' approach to the Ramsey test, namely that conditional sentences express truth-value bearing propositions. Here, we shall point to another assumption that is implicit in Gärdenfors' treatment of conditionals:

> *The Non-Indexicality Assumption*. A conditional sentence A > B expresses *one and the same* proposition relative to every belief state.

We shall argue that — once this assumption is given up — there is no genuine conflict between the Ramsey test and the Preservation condition. That is, it is possible without threat of paradox to keep both the original *Ramsey test*:

(Ramsey) A > B is accepted in a belief state X iff B is accepted in X∗A

and the *Preservation* condition in the form:

(P) If A is consistent with X, then X is included in X∗A,

without giving up the assumption that conditionals express truth-value bearing propositions. The approach that is outlined here is developed in further detail in Lindström (to appear).

When in (P) we say that one belief state X is *included in* another state Y, we mean that all the *propositions* that are accepted in X are also accepted in Y. This does not necessarily mean, however, that all the *sentences* that are accepted in X are accepted in Y. In the presence of context-dependent sentences, that may express different propositions relative to different belief states, inclusion between the propositions accepted does not imply the corresponding inclusion between sentences. Hence, the above form of Preservation does not imply:

> If A is consistent with X, then every sentence that is accepted in X is also accepted in X∗A.

The latter condition is plausible only if all the sentences of the object language are context independent.

5.4.1. The context-sensitive nature of epistemic conditionals

The assumption that the sentences of the object language express determinate propositions in a *context independent* way is implicit in the AGM approach. If one and the same sentence could express different propositions relative to different belief states, then set-theoretic statements concerning belief sets, for instance $G \subseteq H$ or $(A \in G \;\&\; A \in H)$, could not have their intended interpretation. Suppose namely that G and H are belief sets representing the belief states X and Y, respectively. Then, the following condition should hold:

(∗) G is included in H if, and only if, X is included in Y.

If, however, the object language contains context-sensitive sentences, this connection might fail. To see that the right-to-left direction might fail, suppose that every proposition that is accepted in X is also accepted in Y. Let A be a sentence in G and let $[\![A]\!]^X$ be the proposition that A expresses relative to X. Since G represents the state X, $[\![A]\!]^X$ is accepted in X. Then, by the supposition, $[\![A]\!]^X$ is also accepted in Y. But, from this we cannot infer that $A \in H$. For that we would need $[\![A]\!]^Y$ to be accepted in Y which may not be the case, since $[\![A]\!]^X$ and $[\![A]\!]^Y$ may be different propositions. Hence, we cannot conclude that $G \subseteq H$.

To see that also the left-to-right direction of (∗) might fail, suppose that $G \subseteq H$ and that the proposition P is accepted in X. G represents X, so there must be a sentence $A \in G$ such that

$[\![A]\!]^X = P$. Since $G \subseteq H$, $A \in H$. It follows that $[\![A]\!]^Y$ is accepted in Y. We cannot, however, conclude that P is accepted in Y, since $[\![A]\!]^Y$ may be different from P. Once we allow sentences that may express different propositions with respect to different belief states, then both directions of (∗) fail. For context-dependent sentences A, even *Success* fails: A may not be a member of $G*A$.

The approach described here differs from AGM and that of Levi (1988) in making a sharp distinction between the semantic level involving propositions and belief states and the linguistic level involving sentences and sets of sentences. Belief revision is seen as an operation on belief states; and it is primarily propositions rather than sentences that are accepted relative to belief states. We may think of a person's *belief state* as the set of all propositions that he accepts. We do not suppose in general that belief states are logically closed.

It is convenient for our purposes to identify propositions with certain sets of possible worlds. If W is the set of possible worlds, then the set **P** of all the propositions that the agent might entertain is a family of subsets of W. A proposition $P \in \mathbf{P}$ is *true at* a possible world w if, and only if, $w \in P$. We suppose that **P** is a Boolean set algebra, i.e., it contains W and is closed under the Boolean set-operations $\cap$, $\cup$ and -. *Belief states* are certain sets of propositions, i.e., we have a family $\mathbf{K} \subseteq \wp(\mathbf{P})$ of all possible belief states. A proposition P is *accepted* in a belief state X if, and only if, $P \in X$. A belief state X *entails* a proposition P iff $\cap X \subseteq P$. P is *compatible* with X iff $\cap X \cap P \neq \emptyset$. The agent's *theory* T(X) is the set of all propositions that are entailed by his belief state X.

What reasons could we possibly have for saying that conditionals are context-sensitive: that they express different propositions with respect to different belief states? In order to answering this question, let us introduce the notion of a *fallback theory* of X. Intuitively, such a theory is one that may be reached by the agent from his current theory T(X) by deleting propositions that are not "sufficiently" entrenched according to some standard of epistemic entrenchment. To put it differently, a fallback theory of X is a subtheory T of T(X) that is closed upwards under epistemic entrenchment: if $P \in T$ and Q is at least as entrenched as P, then $Q \in T$. In terms of fallback theories, we may give the following interpretation of an epistemic conditional "If A, then B":

A together with some true fallback theory T that is compatible with A entails B.[14]

But what is meant by a fallback theory is dependent on the belief state of the speaker. Given that epistemic conditionals express truth-value bearing propositions, the natural conclusion is that they express different propositions with respect to different belief states: the truth or falsity of an epistemic conditional "If A, then B" is then dependent not only on the world with respect to which the conditional is being evaluated but also on the belief state X of the speaker.[15]

The idea that conditional sentences express different propositions relative to different belief states is quite a natural one. Consider the following two sentences:[16]

(1) If Bizet and Verdi were compatriots, Verdi was French.

(2) If Bizet and Verdi were compatriots, Bizet was Italian.

(1) could be used to make a true statement by a contemporary speaker who knows that Bizet was French, but does not know the nationality of Verdi. For such a speaker, the claim made by (2) would be false. The situation is the opposite for the speaker who knows that Verdi was Italian but does not know the nationality of Bizet.

Instead of assigning propositions to conditional sentences in a context-independent way, we need to *relativize* the assignment of propositions to belief states. Only relative to a belief state does an epistemic conditional A > B express a determinate proposition. We should speak of the proposition $[\![A > B]\!]^X$ expressed by the conditional A > B *relative to* the belief state X. It is then natural to say that the conditional A > B is *accepted* in the belief state X if, and only if, the proposition $[\![A > B]\!]^X$ expressed by A > B relative to X is a member of X. In other words,

A > B is *accepted* in X iff $[\![A > B]\!]^X \in X$.

The analysis of conditionals given here is close to those of Stalnaker (1968) and Lewis (1973), except for containing an *additional parameter:* a belief state. The intuitive idea is expressed by Stalnaker (1975) as follows:[17]

> A conditional statement, *if A, then B*, is an assertion that the consequent is true, not necessarily in the world as it is, but in the world as it would be if the antecedent were true.

In possible worlds terms we can express this idea roughly as follows:

> A conditional sentence A > B is true at a world w just in case B is true at all the A-worlds that are most similar to w.

However, here we shall think of the notion of similarity involved in the truth condition for conditionals as an *epistemic notion* which is determined by the agent's belief state. Making this dependence on a belief state explicit, we get:

> A conditional sentence A > B is *true at a world* w *relative to a belief state* X just in case B is true at all the A-worlds that are most X-similar to w,

where X-similarity is a concept of similarity between possible worlds that is determined by the belief state X. According to this type of semantics, the truth-value of a conditional A > B is dependent both on the state w of the world and the belief state X. Relative to a belief state X, A > B can be said to express the proposition:

$[\![A > B]\!]^X = \{w: A > B \text{ is true at } w \text{ relative to } X\}$.

5.4.2. A solution to the paradox

Once we are reminded of the context dependent nature of conditionals and other epistemic constructions, the representation of belief states by sets of sentences and acceptance by set-theoretic membership in such sets becomes less appealing. If we distinguish between propositions, belief states, acceptance, on the one hand, and sentences, belief sets and membership, on the other, we see that the most perspicuous way of formulating the conditions of Success, Consistency and Preservation is in terms of the former notions:

(P-Success) The proposition P is accepted in $X*P$.

(P-Consistency) If P and X are consistent, when considered separately, then $X*P$ is also consistent.

(P-Preservation) If Q is accepted in a given belief state X and P is consistent with X, then Q is still accepted in $X*P$.

Now, if we formulated the Ramsey test in an analogous fashion as:

(P-R) $P \Rightarrow Q$ is accepted in X iff Q is accepted in $X*P$,

where $\Rightarrow$ is a binary operation on propositions corresponding to the conditional connective >, we would indeed be confronted with Gärdenfors' theorem. We could then derive the following monotonicity condition:

(P-Monotonicity) If $X \subseteq Y$, then $X*P \subseteq Y*P$.

And P-Monotonicity is easily seen to be incompatible with the above conditions on belief revision, given the additional requirement:

(Non-Triviality) There exist two propositions P, Q and three consistent belief states X, Y and Z such that:

(1) $P \in X$ and $X \cup \{-Q\}$ is consistent;

(2) $Q \in Y$ and $Y \cup \{-P\}$ is consistent;

(3) $X \subseteq Z$ and $Y \subseteq Z$.

The proof of this result is a straightforward adaptation of our proof of Gärdenfors' theorem in section 4.

However, thinking of the conditional connective > as corresponding to a binary operation $\Rightarrow$ on propositions is tantamount to assuming that conditional sentences are context-independent. Given such an operation $\Rightarrow$, we could formulate the following semantic clause for conditionals:

(i) $[\![A > B]\!] = [\![A]\!] \Rightarrow [\![B]\!]$,

where $[\![A > B]\!]$ is the proposition expressed by A > B. But if we instead think of conditionals A > B as expressing propositions only relative to belief states, we would rather like to have something like the following semantic clause:

(ii) $[\![A > B]\!]^X = [\![A]\!] \Rightarrow_X [\![B]\!]$,

where $[\![A > B]\!]^X$ is the proposition expressed by A > B relative to the belief state X and $\Rightarrow_X$ is a ternary operation taking two propositions and a belief state as arguments and yielding a proposition as value. Here we assume that the sentences A and B themselves are not context-sensitive, so that the propositions that they express are not dependent on the belief state. For a semantic clause without this restriction, see condition (e) below. We may think of $\Rightarrow_X$ as a *context-dependent operation on propositions.* For such an operation, the Ramsey test takes the form:

(P-Ramsey) $(P \Rightarrow_X Q) \in X$ iff $Q \in X{*}P$.

With P-Ramsey our proof of Gärdenfors' theorem does not go through, since Monotonicity is no longer derivable. To see this suppose that $X \subseteq Y$ and $Q \in X{*}P$. Then, by P-Ramsey, $(P \Rightarrow_X Q) \in X$, from which we conclude $(P \Rightarrow_X Q) \in Y$. However, from this we cannot reach the desired conclusion $Q \in Y{*}P$. To get there we would need $(P \Rightarrow_Y Q) \in Y$, instead.

As a matter of fact, we can prove that there are non-trivial belief revision systems of the type $<W, \mathbf{P}, \mathbf{K}, *, \Rightarrow_X>$ that satisfy the propositional versions of the Gärdenfors axioms for belief revision together with the condition P-Ramsey. That is, we have:

THEOREM. There are systems $\mathcal{S} = <W, \mathbf{P}, \mathbf{K}, *, \Rightarrow_X>$ satisfying Success, Consistency, Preservation, Non-Triviality, P-Ramsey together with the conditions:

(Closure) If X entails P, then $P \in X$;

(W) $X{*}W = X$;

(Revision by Conjunction) If $X{*}P \cup \{Q\}$ is consistent, then $X{*}(P \cap Q) = (X{*}P){+}Q$,

where for any X and P, X+P is the *expansion* of X with P, i.e., the set:

$$\{Q \in \mathbf{P}: \cap X \cap P \subseteq Q\}.$$

Proof: Let W and **P** be given and let **K** be all subsets of **P** that are closed under entailment. We associate with every consistent $X \in \mathbf{K}$ a system $\$_X$ of spheres in the sense of Grove (1988) around $\cap X$. That is, $\$_X$ is a family of subsets of W satisfying the conditions:

(i) $\cap X$ and W belong to $\$_X$;

(ii) for all $S \in \$_X$, $\cap X \subseteq S$;

(iii) for all $S, S' \in \$_X$, $S \subseteq S'$ or $S' \subseteq S$;

(iv) for every $P \in \mathbf{P}$ and every $S \in \$_X$, if $S \cap P \neq \emptyset$, then there exists an S' in $\$_X$ such that $S' \cap P \neq \emptyset$ and for every S'' in $\$_X$, if $S'' \cap P \neq \emptyset$, then $S' \subseteq S''$.

If X and P are consistent, taken separately, then we define $X{*}P$ to be $\{Q \in \mathbf{P}: S \cap P \subseteq Q\}$, where S is the smallest sphere in $\$_X$ such that $S \cap P \neq \emptyset$. Otherwise, we let $X{*}P$ be **P**. It is easily verified that the conditions Success, Consistency, Closure, *(W),* and Revision by Conjunction are satisfied. Preservation follows from *(W)* together with Revision by Conjunction. We can easily see to it that **P** contains two propositions P and Q such that $P \cap Q \neq \emptyset$, $P \cap -Q \neq \emptyset$, $-P \cap Q \neq \emptyset$ and $-P \cap -Q \neq \emptyset$. Two such propositions are said to be completely independent. Let then $X = \{R: P \subseteq R\}$, $Y = \{R: Q \subseteq R\}$ and $Z = \{R: P \cap Q \subseteq R\}$. Then, X, Y, Z are consistent belief states such that:

(1) $P \in X$ and $X \cup \{-Q\}$ is consistent;

(2) $Q \in Y$ and $Y \cup \{-P\}$ is consistent;

(3) $X \subseteq Z$ and $Y \subseteq Z$.

Thus, Non-Triviality is satisfied.

We are next going to define the operation $\Rightarrow_X$. For this purpose, we associate with each world w and each belief state X a system of spheres $\$_{X,w}$ that satisfies the conditions:

(i) W belongs to $\$_{X,w}$;

(ii) for all $S \in \$_{X,w}$, $w \in S$, *(Weak Centering)*

together with the analogues of (iii) and (iv) for $\$_{X,w}$. We then impose the following constraint:

if $w \in \cap X$, then $\$_{X,w} = \$_X$. *(Compatibility)*

That is, if w is a world that is compatible with all the beliefs in state X, then the sphere system around w coincides with that around X.

We define $\Rightarrow_X$ by letting:

$$P \Rightarrow_X Q = \{w: (\exists S \in \$_{X,w})(\varnothing \neq S \cap P \subseteq Q)\}.$$

It remains to show that P-Ramsey holds, i.e.,

$$P \Rightarrow_X Q \in X \text{ iff } Q \in X{*}P.$$

Suppose that $P \Rightarrow_X Q \in X$. Then, $\cap X \subseteq P \Rightarrow_X Q$. That is, (i) for all $w \in \cap X$, $w \in (P \Rightarrow_X Q)$. If $\cap X = \varnothing$, then $X{*}P = \mathbf{P}$. Hence, the desired conclusion holds in this case. Suppose that (ii) $\cap X \neq \varnothing$. By the constraint: for all $w \in \cap X$,

$$w \in (P \Rightarrow_X Q) \text{ iff } (\exists S \in \$_X)(\varnothing \neq S \cap P \subseteq Q).$$

But (i) and (ii) yields that for some $w \in \cap X$, $w \in (P \Rightarrow_X Q)$. Hence, $(\exists S \in \$_X)(\varnothing \neq S \cap P \subseteq Q)$. But this means that $Q \in X{*}P$.

For the other direction, suppose that $Q \in X{*}P$. Let $w \in \cap X$. By the constraint,

$$w \in (P \Rightarrow_X Q) \text{ iff } (\exists S \in \$_X)(\varnothing \neq S \cap P \subseteq Q).$$

That is,

$$w \in (P \Rightarrow_X Q) \text{ iff } Q \in X{*}P.$$

It follows that $w \in (P \Rightarrow_X Q)$. We have shown that, $\cap X \subseteq (P \Rightarrow_X Q)$, which means that $(P \Rightarrow_X Q) \in X$. □

In the above proof, we outlined a semantics for belief revision and conditionals based on systems of spheres. First, every belief state X was associated with a system of spheres $\$_X$ in terms of which the belief revision operation X∗... was defined. Secondly, each world w was associated with a system of spheres $\$_{X,w}$ relative to X. In terms of the latter system of spheres, we could define the propositional operation (... $\Rightarrow_X$ ...). A condition was imposed connecting the two kinds of sphere systems:[18]

if w is compatible with all the beliefs in X, then $\$_{X,w} = \$_X$.

From this condition, we proved:

(P-Ramsey) $P \Rightarrow_X Q \in X$ iff $Q \in X*P$.

This modelling showed (P-Ramsey) to be compatible with propositional versions of Gärdenfors' axioms for belief revision.

Suppose now that we have a formal language $\mathcal{L}$ with sentences built up from atomic ones using Boolean connectives $\bot$ and $\rightarrow$ and the conditional connective $>$. $\mathcal{L}_0$ is the fragment of $\mathcal{L}$ without $>$. We let $\mathcal{S} = \langle W, \mathbf{P}, \mathbf{K}, *, \Rightarrow_X\rangle$ be a belief revision system satisfying Success, Consistency, Preservation, Non-Triviality, Closure and P-Ramsey. We let $[\![...]\!]$ be an interpretation function that assigns propositions $[\![A]\!]^X$ to sentences A in $\mathcal{L}$ relative to belief states X. This function is assumed to satisfy the requirements:

(a) If A is a sentence of $\mathcal{L}_0$, then for all $X, Y \in \mathbf{K}$, $[\![A]\!]^X = [\![A]\!]^Y$. Hence, for sentences of $\mathcal{L}_0$ we may write $[\![A]\!]$ instead of $[\![A]\!]^X$.

(b) For every $P \in \mathbf{P}$, there exists a sentence A of $\mathcal{L}_0$ such that $P = [\![A]\!]^X$.

(The expressibility assumption)

(c) $[\![\bot]\!]^X = \varnothing$.

(d) $[\![A \rightarrow B]\!]^X = (W - [\![A]\!]^X) \cup [\![B]\!]^X$.

(e) $[\![A > B]\!]^X = [\![A]\!]^X \Rightarrow_X [\![B]\!]^{(X*[\![A]\!]^X)}$

Writing $X*A$ for $X*[\![A]\!]^X$, we can simplify (e) to:

$$[\![A > B]\!]^X = [\![A]\!]^X \Rightarrow_X [\![B]\!]^{(X*A)}.$$

Assumption (a) says that the sentences of the basic language $\mathcal{L}_0$ are context-independent. The expressibility assumption is the requirement that the basic language has sufficient expressive power to express all the propositions that the agent might accept. Together these two conditions makes it possible to represent belief states, in a context-independent way, by sets of sentences of $\mathcal{L}_0$.

We say that a sentence A is *accepted* in the belief state X just in case the proposition $[\![A]\!]^X$ that is expressed by A relative to X is accepted in X. That is:

A is accepted in X iff $[\![A]\!]^X \in X$.

Then, we get:

A > B is accepted in X iff
$[\![A > B]\!]^X \in X$ iff
$([\![A]\!]^X \Rightarrow_X [\![B]\!]^{(X*A)}) \in X$ iff
$[\![B]\!]^{(X*A)} \in X*[\![A]\!]^X$ iff
$[\![B]\!]^{(X*A)} \in X*A$ iff
B is accepted in $X*A$.

That is, we get our original formulation of the Ramsey test:

(Ramsey) A > B is accepted in the belief state X iff B is accepted in $X*A$.

5.4.3. Belief revision at the linguistic level

Our resolution of Gärdenfors' paradox depended on viewing belief revision as primarily an operation on belief states and interpreting Gärdenfors' postulates on belief revision as applying to such an operation. We showed that such a propositional belief revision system could be provided with a context-dependent operation $\Rightarrow_X$ on propositions satisfying (P-Ramsey). Finally, we showed that the conditional connective > could be interpreted semantically in terms of $\Rightarrow_X$ in such a way that the Ramsey test became valid.

Now, we want to see what happens when we view belief revision as an operation on belief sets, i.e., sets of sentences, instead. Starting out from a belief revision system $\mathcal{S} = \langle W, \mathbf{P}, \mathbf{K}, *, \Rightarrow_X \rangle$ and an interpretation function $[\![\ldots]\!]$ satisfying the conditions (a) - (e) above, we define a corresponding logic L, a set $\uparrow(\mathbf{K})$ of belief sets corresponding to the set **K** of all belief states, and an operation $*$ of belief revision on belief sets. As a matter of fact, we define two notions of belief set, one for the basic language $\mathcal{L}_0$ and one for the extended language $\mathcal{L}$, and correspondingly two notions of belief revision. Within our framework, the two notions are interdefinable and the Ramsey test can be formulated in terms of both.

First, we define the *logic* L determined by $\mathcal{S}$ and $[\![\ldots]\!]$. We say that a sentence A in $\mathcal{L}$ is an L-*consequence* of a set Γ of sentences in $\mathcal{L}$ (in symbols, $\Gamma \vdash_L A$) if for every belief state X $\in$ **K**, $\bigcap\{[\![B]\!]^X : B \in \Gamma\} \subseteq [\![A]\!]^X$. That is, $\Gamma \vdash_L A$ iff for every belief state X and every possible world w, if all the sentences in Γ are true at w relative to X, then A is also true at w relative to X. For sentences in $\mathcal{L}_0$ the reference in this definition to the belief state X becomes superfluous. That is, if Γ is a set of sentences in $\mathcal{L}_0$ and A belongs to $\mathcal{L}_0$, then: $\Gamma \vdash_L A$ iff $\bigcap\{[\![B]\!] : B \in \Gamma\} \subseteq [\![A]\!]$.

Next, we need to decide on what we shall understand by a belief set. Each belief state is associated with two sets of sentences: First we have the set:

$$\{A \in \mathcal{L}_0 : [\![A]\!] \in X\}$$

of all *non-indexical* or *basic sentences* that correspond to propositions in X. Then there is the set of all sentences of the extended language $\mathcal{L}$ that are accepted in X, i.e., the set:

$$\{A \in \mathcal{L} : [\![A]\!]^X \in X\}.$$

Let us speak of the first set as the *descriptive belief set* corresponding to the belief state X, and the second set as the *acceptance set* corresponding to X. In view of the expressibility assumption, there is a one-to-one correspondence between belief states and descriptive belief sets. We also have a one-to-one correspondence between descriptive belief sets and acceptance sets. For each acceptance set G, the corresponding descriptive belief set is the set G $\cap$ $\mathcal{L}_0$ which we may refer to as the *descriptive core* of G, or core(G). Conversely, for each descriptive belief set K, we can define the corresponding acceptance set as:

$$E(K) = \{A \in \mathcal{L} : [\![A]\!]^{[\![K]\!]} \in [\![K]\!]),$$

where $[\![K]\!]$ is the belief state that corresponds to K, i.e., $[\![K]\!] = \{[\![B]\!]: B \in K\}$. In other words, E(K) is the set of all sentences of $\mathcal{L}$ that are accepted in the belief state corresponding to K. Of course, G = E(K) iff K = core(G).

For each belief state X, we let $\uparrow(X)$ be the acceptance set corresponding to X, i.e.,

$$\uparrow(X) = \{A \in \mathcal{L}: [\![A]\!]^X \in X\}.$$

and we let $\uparrow(\mathbf{K})$ be the set of all acceptance sets, i.e.,

$$\uparrow(\mathbf{K}) = \{\uparrow(X): X \in \mathbf{K}\}.$$

The sentences of core(G) are context-independent, so we can speak in a context-independent way of the proposition $[\![A]\!]$ expressed by A, for each $A \in$ core(G). Furthermore, we have assumed that every proposition in **K** is expressed by some sentence in $\mathcal{L}_0$ (The Expressibility Assumption). It follows that we can recover the belief state corresponding to a acceptance set G as the set of all propositions that are expressed by some member of core(G). That is, the belief state corresponding to G is defined as:

$$[\![G]\!] = \{[\![A]\!]: A \in \text{core}(G)\} = [\![\text{core}(G)]\!].$$

Notice that:

$$[\![\uparrow(X)]\!] = X, \text{ and}$$

$$\uparrow([\![G]\!]) = \{A \in \mathcal{L}: [\![A]\!]^{[\![G]\!]} \in [\![G]\!]\} = \{A \in \mathcal{L}: A \in G\} = G.$$

We also have:

$$\text{if } [\![G]\!] = [\![H]\!], \text{ then } G = H.$$

In order to prove this, let $[\![G]\!] = [\![H]\!]$. Then, $G = \uparrow([\![G]\!]) = \uparrow([\![H]\!]) = H$.

We can now define two operations of belief revision, one operation $\oplus$ on descriptive belief sets and the other $*$ on acceptance sets. For any descriptive belief set K and any $A \in \mathcal{L}$, we let:

$$K \oplus A = \{B \in \mathcal{L}_0: [\![B]\!] \in [\![K]\!]*A\} = \{B \in \mathcal{L}_0: [\![B]\!] \in [\![K]\!]*[\![A]\!]^{[\![K]\!]}\}.$$

That is, if K is a descriptive belief set and A is a sentence of $\mathcal{L}$, then we define $K \oplus A$ as follows: first, we go to the belief state $[\![K]\!]$ corresponding to K. We then revise that state with the proposition $[\![A]\!]^{[\![K]\!]}$ that A expresses relative to that state. Finally, we let $K \oplus A$ be the set of all sentences of $\mathcal{L}_0$ that are accepted in the resulting belief state.

Similarly, we define for any acceptance set G:

$$G*A = \uparrow([\![G]\!]*A) = \{B \in \mathcal{L}: [\![B]\!]^{([\![G]\!]*A)} \in [\![G]\!]*A\}.$$

We have:

$$[\![K \oplus A]\!] = [\![K]\!]*A, \text{ and}$$

$$[\![G*A]\!] = [\![G]\!]*A.$$

The two operations are interdefinable as follows:

$G*A = E(core(G) \oplus A)$
$K \oplus A = core(E(K)*A)$.

For any pair of descriptive belief states K, K' we have:

$K \subseteq K'$ iff $[\![K]\!] \subseteq [\![K']\!]$.

However, for acceptance sets G, H, we do *not* have:

$G \subseteq H$ iff $[\![G]\!] \subseteq [\![H]\!]$.

For acceptance sets G and H, it is important not to conflate ordinary set inclusion ($G \subseteq H$) with the relation (we write it, $G \sqsubseteq H$) that holds iff all the *propositions* that are accepted in the belief state $[\![G]\!]$ are also accepted in $[\![H]\!]$. Due to the Expressibility Assumption, we can define $\sqsubseteq$ as follows:

$G \sqsubseteq H$ iff $core(G) \subseteq core(H)$.

We have, of course,

$G \sqsubseteq H$ iff $[\![G]\!] \subseteq [\![H]\!]$.

Let us now see how to formulate the Ramsey test and Preservation within the present framework. First, we consider the Ramsey test:

(Ramsey) A > B is accepted in the belief state X iff B is accepted in X∗A.

In terms of acceptance sets and revision of acceptance sets, this becomes:

$A > B \in G$ iff $B \in G*A$.

The same condition formulated in terms of descriptive belief sets K and the operation $\oplus$ is:

$A > B \in E(K)$ iff $B \in E(K \oplus A)$.

Consider next P-Preservation:

If Q is accepted in a given belief state X and P is consistent with X, then Q is still accepted in X∗P. *(P-Preservation)*

This corresponds to:

If $A \in \mathcal{L}_0$, K is a descriptive belief set and $K \cup \{A\} \nvdash_L \bot$, then $K \subseteq K \oplus A$, *($\mathcal{L}_0$-Preservation)*

In other words,

If $A \in \mathcal{L}_0$, G is an acceptance set and $(core(G) \cup \{A\}) \nvdash_L \bot$, then $G \sqsubseteq G*A$.

Suppose next that $\mathcal{S} = \langle W, \mathbf{P}, \mathbf{K}, *, \Rightarrow_X \rangle$ satisfies P-Success, P-Consistency, P-Preservation, Closure, (W), Revision by Conjunction, Non-Triviality and P-Ramsey and that $[\![...]\!]$ satisfies the conditions (a) - (e) above. Then, the following conditions are also satisfied:

(1) The logic L determined by $\mathcal{S}$ and $[\![...]\!]$ contains all substitution instances of truth-functional tautologies and is closed under modus ponens (i.e., if $\vdash_L A$ and $\vdash_L A \rightarrow B$, then $\vdash_L B$).

(2) If $\vdash_L A \leftrightarrow B$ and $\vdash_L C \leftrightarrow D$, then $\vdash_L (A > C) \leftrightarrow (B > D)$.

(3) Descriptive belief sets and acceptance sets are L-closed sets in $\mathcal{L}_0$ and $\mathcal{L}$, respectively. *(Closure)*

(4) E is a one-to-one mapping between descriptive belief sets and acceptance sets such that for each descriptive belief set K, $K = E(K) \cap \mathcal{L}_0$.

(5) If $A \in \mathcal{L}_0$, then $A \in K \oplus A$. *($\mathcal{L}_0$-Success)*

(6) If $A \nvdash_L \bot$ and $K \nvdash_L \bot$, then $K \oplus A \nvdash_L \bot$. *($\mathcal{L}_0$-Consistency)*

(7) If $A \in \mathcal{L}_0$, K is a descriptive belief set and $K \cup \{A\} \nvdash_L \bot$, then $K \subseteq K \oplus A$, *($\mathcal{L}_0$-Preservation)*

(8) If $A \in \mathcal{L}_0$ and $K \cup \{A\} \nvdash_L \bot$, then $K \oplus A = K + A$,
where $K+A = \{B \in \mathcal{L}_0 : K \cup \{A\} \vdash_L B\}$. *($\mathcal{L}_0$-Expansion)*

(9) If $\vdash_L A \leftrightarrow B$, then $K \oplus A = K \oplus B$. *(Substitutivity of logical equivalents)*

(10) If $A, B \in \mathcal{L}_0$ and $K \oplus A \cup \{B\} \nvdash_L \bot$, then $K \oplus (A \wedge B) = (K \oplus A) + B$. *($\mathcal{L}_0$-Revision by Conjunction)*

(11) $A > B \in E(K)$ iff $B \in E(K \oplus A)$. *(RT)*

(12) There exist two sentences B and C in $\mathcal{L}_0$ and three consistent descriptive belief sets G, H and K such that: (i) $B \in G$ and $G \cup \{\neg C\}$ is consistent; (ii) $C \in H$ and $H \cup \{\neg B\}$ is consistent; and (iii) $G \subseteq K$ and $H \subseteq K$. *($\mathcal{L}_0$-Non-Triviality)*

In virtue of Theorem 2, there are belief revision systems satisfying the above conditions. It is also easy to see that $[\![...]\!]$ can be defined (recursively) in such a way that conditions (a) - (e) are satisfied. It follows that conditions (1) - (12) are mutually consistent.

The present approach has the formal advantage over Levi's (1988) of being able to account for iterated conditionals in a natural way. Levi's version of the Ramsey test does not provide a method for evaluating such conditionals. The present version of the test can, however, be applied to iterated conditionals without difficulty. Consider, for example, $(A > B) > (C > D)$. According to (RT), we have:

$$(A > B) > (C > D) \in E(K) \text{ iff } C > D \in E(K \oplus (A > B)) \text{ iff } D \in E((K \oplus (A > B)) \oplus C).$$

Or, in other words,

$$(A > B) > (C > D) \in G \text{ iff } C > D \in G*(A > B) \text{ iff } D \in (G*(A > B))*C.$$

Semantically this means:

$$[\![(A > B) > (C > D)]\!]^X \in X \text{ iff } [\![C > D]\!]^{X*(A > B)} \in X*(A > B) \text{ iff } [\![D]\!]^{(X*(A > B))*C} \in (X*(A > B))*C,$$

where X is the belief state $[\![G]\!]$.

5.4.4. Could the paradox be reinstated?

It should be pointed out that the following form of Monotonicity:

For any descriptive belief sets K, K', and any $A \in \mathcal{L}_0$, if $K \subseteq K'$, then $K \oplus A \subseteq K' \oplus A$. ($\mathcal{L}_0$-Monotonicity)

is sufficient in the presence of $\mathcal{L}_0$-Success, $\mathcal{L}_0$-Consistency, $\mathcal{L}_0$-Preservation and $\mathcal{L}_0$-Non-Triviality for the derivation of an inconsistency. However, it is impossible to derive $\mathcal{L}_0$-Monotonicity from:

(RT) For every descriptive belief set K and any $A, B \in \mathcal{L}$,
$A > B \in E(K)$ iff $B \in E(K \oplus A)$.

Thus, Gärdenfors' paradox is avoided.

At this point the reader might object and point out that there is another form of Monotonicity that actually follows from (RT), namely:

For any descriptive belief sets K, K', and any $A \in \mathcal{L}$,
if $E(K) \subseteq E(K')$, then $E(K \oplus A) \subseteq E(K' \oplus A)$. ($\mathcal{L}$-Monotonicity)

Couldn't this condition be used to construct a version of Gärdenfors' paradox? This is in fact possible. The only thing we have to do is to replace the condition of $\mathcal{L}_0$-Non-Triviality with the stronger condition:

There exist two sentences B and C in $\mathcal{L}_0$ and three consistent descriptive belief sets G, H and K such that: (i) $B \in G$ and $G \cup \{\neg C\}$ is consistent; (ii) $C \in H$ and $H \cup \{\neg B\}$ is consistent; and (iii') $E(G) \subseteq E(K)$ and $E(H) \subseteq E(K)$. *($\mathcal{L}$-Non-Triviality)*

It is easy to see that this condition is sufficient to derive a contradiction. Let A be the sentence $\neg B \vee \neg C$. Since, B and C belong to $\mathcal{L}_0$, the same holds for A. It follows from (i) and (ii) that each of G and H are logically compatible with A. Since, B and C belong to G and H, respectively, $\mathcal{L}_0$-Preservation implies that $B \in G \oplus A$ and $C \in H \oplus A$. Hence, $B \in E(G \oplus A)$ and $C \in E(H \oplus A)$ (since, for any descriptive belief sets G, $E(G) \cap \mathcal{L}_0 = G$). However, since $E(G)$, $E(H) \subseteq E(K)$ (condition (iii')), $\mathcal{L}$-Monotonicity implies that $E(G \oplus A)$, $E(H \oplus A) \subseteq E(K \oplus A)$. It follows that $B, C \in E(K \oplus A)$. By $\mathcal{L}_0$-Success, we also get $A \in E(K \oplus A)$. But this implies that $K \oplus A$ is inconsistent. On the other hand, $K \oplus A$ must be consistent, by $\mathcal{L}_0$-Consistency.

It might seem as if we have indeed succeeded in reinstating the paradox. However, this is not really so. The above proof is nothing but a reductio proof of the negation of $\mathcal{L}$-Non-Triviality from the premises: $\mathcal{L}_0$-Success, $\mathcal{L}_0$-Consistency, $\mathcal{L}_0$-Preservation and (RT). Given these assumptions, the situation envisaged in $\mathcal{L}$-Non-Triviality is impossible. This is no paradox, since we have no reasons to believe $\mathcal{L}$-Non-Triviality to be true. The situation here is completely analogous to the one we encountered before in connection with the attempt to prove Gärdenfors' theorem for Levi's version of the Ramsey test. $\mathcal{L}$-Non-Triviality might appear reasonable if we do not distinguish clearly between descriptive belief sets and acceptance sets or if we conflate inclusion between descriptive belief sets with the same relation among acceptance sets.

6. Summing up

We have seen that Gärdenfors' paradox is based on a number of questionable assumptions and that it can be resolved in various ways. In particular, we have discussed three major approaches to the paradox. According to the first approach we divided the Ramsey test into two logically independent conditions: Strict RT and Monotonicity. The proposal was to replace the Ramsey test by Strict RT while abandoning Monotonicity. In the absence of Monotonicity, it is possible to distinguish between the conditions:

(1) $B \in G{*}A$; and

(2) for every extension H of G, $B \in G{*}A$,

while Monotonicity is just the assumption that they are equivalent. The idea behind Strict RT is to demand the logically stronger of these conditions, namely (2), for the conditional A > B to be a member of the belief set G. According to this analysis, the fault with the original Ramsey test is that it leads to the collapse of the intuitively distinct conditions (1) and (2). Gärdenfors' paradox is avoided, since it essentially involves the assumption of monotonicity.

The second approach that we considered, Levi's non-propositional one, also involved modifying the Ramsey test. However, here the intuitive idea was different: rather than expressing truth-value bearing propositions, epistemic conditionals express policies for the revision of belief states. This leads to a modified Ramsey-test according to which an epistemic conditional A > B, involving ordinary descriptive sentences A and B, is accepted in a belief state G if and only if $B \in G{*}A$. Since, conditionals cannot themselves be members of belief states Monotonicity is not derivable so Gärdenfors' paradox is avoided.

Finally, we considered an approach, the indexical one, according to which epistemic conditionals were thought of as expressing genuine beliefs, but only in a context-sensitive manner, relative to a belief state. A distinction was made between an acceptance set, the set of all sentences including context-sensitive ones that are accepted in a belief state, and its descriptive core (a descriptive belief set) consisting of all the non-indexical sentences that are accepted in a belief state. We presented a semantics that validates both the full Ramsey test at the level of acceptance sets and Gärdenfors' axioms for belief revision, provided that the latter are applied to descriptive belief sets and not to acceptance sets. By not conflating the level of descriptive belief sets with that of acceptance sets Gärdenfors' paradox is avoided.

The three different ways of resolving the paradox need not be competitors to each other. Instead they might be viewed as corresponding to different uses of conditional constructions in epistemic contexts. However, one question still remains: Is the Ramsey test intuitively plausible as a claim about our everyday use of conditionals? This is the question that we now want to address.

But first we must say something about how we want to interpret the notion of belief. Here, we prefer to follow the pragmatist tradition. In the words of Peirce: "*belief* consists mainly in being deliberately prepared to adopt the formula believed in as a guide to action".[19] Thus, beliefs are to be interpreted as what we have previously called "assumptions".

As we remember, revision of assumption sets violates Preservation. This is why we could not take this interpretation of belief for granted while exploring different ways of dealing with Gärdenfors' paradox. Now, when the question concerns the intuitive plausibility of the Ramsey test itself, we are free to opt for the interpretation of belief that we find most attractive.

Let us first consider the 'if'-part of the Ramsey test *(RT⇐)* . That this part of the test is quite counterintuitive has already been pointed out by Gärdenfors in his book:

> The most problematic implication of (RT) is the one saying that if $B \in G{*}A$, then $A > B \in G$. In a sense, this implication requires that too many conditionals be elements of a belief set G because it contains conditionals related to all possible revisions that G may undergo.[20]

In fact, as argued by Nils-Eric Sahlin[21], Ramsey himself would probably reject that part of the test that bears his name.

We give an example to indicate that *(RT⇐)* generates "too many" conditionals.[22] Suppose that Oscar, in his present state of belief (represented by the set) G believes that Tweety is a bird, that Tweety has been reported to fly by a normally reliable witness, that normally birds can fly and that penguins are birds. However, he has no opinion concerning penguins ability to fly. Let A be the proposition that Tweety is a penguin and B the proposition that Tweety can fly. In view of his belief that it is normal for birds to fly and that Tweety actually has been reported to fly, Oscar believes B. He is prepared to act on that assumption. On the other hand, he does not know whether A is true or not.

Now, if Oscar were to learn (A) that Tweety is a penguin, he would still believe (B) that Tweety can fly. That is,

(1) $B \in G{*}A$.

Thus, if *(RT⇐)* were a valid principle, we would have:

(2) $A > B \in G$,

that is, Oscar would already in his present state G believe $A > B$.

Consider now the proposition (C) that penguins cannot fly. As we have pointed out, C is compatible with Oscar's beliefs in G. Hence, the conditional $A > \neg B$, which intuitively is entailed by C in view of the law-like character of that proposition, must also be compatible with Oscar's beliefs in G. But this contradicts (2). The argument, of course, depends on the following intuitive principle:

> Conditionals with the same logically possible antecedent and with incompatible consequents are mutually incompatible.

We conclude from this argument, that *(RT⇐)* should be rejected. (At least for the assumption-interpretation of beliefs. Were we to interpret beliefs as certainties, the example would not work: given his original evidence, it would be unreasonable for Oscar to be certain of

Tweety's ability to fly. Nor would the example work for the expectation-interpretation. Oscar who originally believes that birds normally can fly, expects that the same applies to penguins. Thus, C and the conditional that C entails, $A > \neg B$, conflict with Oscar's original expectations.)

What about the other direction of the test, its 'only if'-part?

Gärdenfors (1988, p 166) presents an example which is supposed to undermine (RT⇒). This example, however, is easy to dismiss: the conditional which figures in it is not epistemic but *ontic.*[23] And we already know that ontic conditionals fail to satisfy the 'only if'-part of the Ramsey test: I believe that no one would have killed Kennedy in Dallas if Oswald had not done it, but I would, of course, come to accept the theory about another murderer, if I were to learn that Oswald in fact was innocent.

However, in connection with this unhappy example, Gärdenfors makes some general observations which are less easy to ignore. According to the Consistency Postulate, if G and the new information A are internally consistent, $G{*}A$ is a consistent set of beliefs even when A conflicts with the original belief set G. This means that revising G with A demands that we give up some of our original beliefs in order to "make room" for A. We should try to keep as many of our old beliefs as possible intact; otherwise we would be throwing away lots of babies with the bath water. But some of the original beliefs must be given up, if the new information A is to be made consistent with the rest of our beliefs.

To begin with, we have to give up our original belief in non-A. This is clear. But normally we have to make other adjustments as well. Thus, suppose that we originally accept two propositions, C and C', which are such that A is compatible with each of them taken by itself but entails that at least one of them must be false. Thus, we have to remove at least one of these propositions when we learn that A is true. If we do not want to remove both, which one should we give up? Gärdenfors suggests giving up the one that is less *entrenched* in our original belief set — the one that is less useful to us "in inquiry and deliberation".

> The fundamental criterion for determining the epistemic entrenchment of a sentence is how useful it is in inquiry and deliberation. Certain pieces of our knowledge and beliefs about the world are more important than others when planning future actions, conducting scientific investigations, or reasoning in general. To give an example of the scientific case, in modern chemical theory, knowledge about combining weights is much more important for chemical experiments than knowledge about the color or taste of some substances. This difference in entrenchment is reflected in that, if chemists for some reason changed their opinion concerning the combining of weights of two substances, this would have much more radical effects on chemical theory than if they changed their opinion concerning the tastes of the two substances. (*Ibid.*, p. 87)

For pragmatical reasons, beliefs that are more entrenched in this sense are more immune to revision, less vulnerable to removal from the original set of beliefs. We keep them if we can.

Now, suppose that I originally believe that non-A and non-B, but that *if* A, *then* B. That is, I accept the conditional $A > B$. Suppose that I then learn that, contrary to what I have believed, A is true. To make room for this new information I must give up one of my old beliefs: either non-B or $A > B$. I cannot cling to both of them, since they are together inconsistent with the new information A. (Here, we assume that epistemic conditionals obey Modus Ponens: From A > B and A, one can derive B.) I know that I should give up that belief which is less entrenched. But according to the 'only if'-part of the Ramsey test, if A > B has belonged to the original belief set and I revise that set with A, then B should always belong to the revised set. Thus, it is non-B that should always be given up: non-B is never more entrenched than A > B if (RT⇒) is generally valid. Is it reasonable? Is it reasonable to assume that epistemic conditionals are always so well entrenched in our beliefs (as compared to the negations of their consequents)? Gärdenfors doubts it. (*Ibid.*, p. 166)

Once one starts to doubt, the examples of epistemic conditionals that seem to violate (RT⇒) are easy to find. To take an extreme case, consider the following exchange:

Me: Oswald did it!

The devil's advocate: Are you sure?

Me: How can you doubt it given all the evidence? If Oswald didn't do it, then I am the emperor of China!

Lucifer: I hate to interrupt your interesting discussion, but, as a matter of fact, Oswald was innocent.

Me: What? Oh, I see. Thank you for putting me straight.

The devil's advocate: Are you then the emperor of China?

Me: Save your jokes for another occasion, will you?

The conditional "If Oswald did not do it, then I am the emperor of China", which I apparently accept before Lucifer's intervention, violates (RT⇒): upon learning the antecedent of this conditional, I am not at all prepared to accept its consequent. The negation of the consequent is much too well entrenched in my original beliefs to be given up.

It is not quite clear to us how convincing such examples are. The defender of Ramsey might say that his test is meant to apply only to rational persons and only to conditionals which such persons "really" accept. Do I really accept that I must be the emperor of China if I am wrong in my belief in Oswald's guilt? Or is it only a hyperbole, an exaggeration used by me to make my point more strongly than I really am entitled to? In asserting the conditional, I want to suggest that my belief in Oswald's guilt is extremely well entrenched — that its degree of entrenchment is comparable to my conviction that I am not the emperor of China. Of course, this is overstating things quite a bit. Being a reasonable person, I do not really accept the conditional in question, as shown by my behavior after Lucifer has volunteered the new information.

This is how the defender of Ramsey might respond. That epistemic conditionals, if sincerely believed by rational people, are always more entrenched than the negations of their antecedents is simply a reflection of the essential connection obtaining between such conditionals and belief change: the connection that the Ramsey test tries to articulate.[24]

In fact, this response begs the question somewhat, as witnessed by the last sentence in the next-to-last paragraph. The claim made in this sentence will be convincing only to someone who already accepts the 'only if'-part of the Ramsey test.

The authors of the present paper — do. That is, we are convinced. What about you, our reader?

NOTES

[*] The present chapter (or rather, a close predecessor of it) has appeared in *Theoria* 58 (1992). Parts of it have grown out of Rabinowicz (1991), which is a non-technical discussion of the approach to the Ramsey test that we develop, in full technical detail, in Lindström and Rabinowicz (1992). Section 5.4 on the "indexical" approach is based on Lindström (to appear). The ideas of this chapter have also been presented at workshops and conferences in Lund, Konstanz, Uppsala, Aix en Provence and Gif-sur-Yvette. We are indebted to the participants, especially André Fuhrmann, Peter Gärdenfors, Sven-Ove Hansson, David Makinson, Michael Morreau, Hans Rott and Krister Segerberg, for comments and suggestions. We are also grateful to Isaac Levi for his thought-provoking written comments.

[1] Department of Philosophy and Philosophy of Science, Umeå University, S-90187 Umeå, Sweden and Department of Philosophy, Uppsala University, S-752 36 Uppsala, Sweden.

[2] For his philosophical production, see *F. P. Ramsey's Philosophical Papers,* edited by D. H. Mellor, Cambridge University Press, Cambridge 1990.

[3] The quotation comes from a paper entitled 'General propositions and causality', written in 1929 (see Ramsey (1931), p.248) (our italics).

[4] Stalnaker (1968). For Stalnaker, this test represents the acceptability condition ("belief condition") for conditionals, which he distinguishes from their *truth condition*. In Lindström and Rabinowicz (1992), we make a similar distinction, even though we ultimately modify the Ramsey test for acceptability and propose a truth condition that differs from Stalnaker's. See Section 5.2 below.

[5] Dudman (1984), (1988), (1991).

[6] Dudman (1984), pp. 152, 153.

[7] ibid., p. 153.

[8] The resemblance should not be overstated. Unlike ourselves, Dudman emphatically claims that neither conditionals nor hypotheticals express any propositions. The former express our attitudes towards our own projective fantasies, while the latter may be seen as compressed arguments. Dudman also rejects the Ramsey test, at least as this test is normally formulated, mainly because he takes the arguments expressed in hypotheticals to be dependent on our particular commitments in a given context and not on the beliefs that we happen to have (Cf. Dudman, 1991, 228-229).

[9] The underlying logic is assumed to satisfy the following requirements: (i) classicality (it extends classical propositional logic); (ii) closure under modus ponens; (iii) the deduction theorem; (iv) compactness. As a

consequence, it also satisfies Lindenbaum's Lemma which says that every consistent set of sentences is included in a maximally consistent set.

[10] Here we have chosen a formulation of the Ramsey test which implies that G∗A is inconsistent whenever G is. This consequence can be avoided by restricting the test to consistent belief sets G only. Our reasoning would not be affected if we modified the Ramsey test in this way and made the corresponding changes in the other conditions that we consider.

[11] As we shall argue in section 5.4, condition (3) in the Non-Triviality condition, as it is stated here, is in fact too strong in the presence of the Ramsey test. The present formulation of Non-Triviality corresponds to what we refer to there as $\mathcal{L}$-Non-Triviality. We argue that this condition should be abandoned and replaced by the weaker condition $\mathcal{L}_0$-Non-Triviality. With this change Gärdenfors' proof does not go through.

[12] Through all of Section 5.2 we assume that every logically closed set of formulas is a possible belief set. This assumption is used in the proof of Weak RT from Normality. Suppose that Normality holds. The left-to-right direction of Weak RT follows immediately. So we prove the direction from right to left. Suppose that B ∈ H∗A for every opinionated extension H of G. Normality then yields that A > B ∈ H, for every opinionated extension H of G. By the assumption that all logically closed sets of formulas are possible belief sets, we obtain that A > B ∈ H for all maximally consistent sets of sentences that extend G. By Lindenbaum's Lemma, this entails that A > B is a logical consequence of G. But G is logically closed, so A > B ∈ G.

[13] To solve this problem Hansson (1992) introduces a primitive three-place relation S(X, Y, Z) between belief states with the intuitive meaning: X is at least as similar to Y as is Z. Hansson then proposes the following *similarity-based Ramsey test:* A > B is accepted in X iff B is accepted in all the S-closest states to X in which A is accepted. He proves that Levi's Ramsey test can be seen as a special case of the similarity-based test. Hansson's condition, of course, allows for the iteration of conditionals. Another proposal for solving the iteration problem has been outlined by Isaac Levi himself in private communication.

[14] Here, we are assuming, for the sake of simplicity, that A and B do not themselves contain conditionals.

[15] In reality, the belief state against which an epistemic conditional is evaluated may not be the agent's actual belief state but rather some hypothetical belief state that is provided by context. It may for example consist of the shared beliefs among the participants in some discussion. This point was emphasised by Isaac Levi in private communication.

[16] Cf. Quine (1962), p. 15.

[17] Stalnaker (1975). See Jackson (1991), p. 143.

[18] Given $\$_X$, we could, for each w ∈ W, *define* $\$_{X,w}$ as the set of spheres S ∈ $\$_X$ such that w ∈ S. It is easily seen that $\$_{X,w}$, when defined in this way, satisfies the conditions that we formulated for a system of spheres around a world w and that Compatibility is also satisfied. Hence, by means of this construction, we have shown that it is possible to satisfy Compatibility and, consequently also P-Ramsey. In addition, we get the following semantic clause for conditional propositions:

$$\text{for all } w \in W,\ w \in P \Rightarrow_X Q \text{ iff } (\exists S \in \$_X)(w \in S,\ S \cap P \neq \emptyset \text{ and } S \cap P \subseteq Q).$$

That is, the proposition $P \Rightarrow_X Q$ is *true at* w iff there is a sphere S around ∩X such that w ∈ S, S is P-permitting, and S ∩ P entails Q. Intuitively, the spheres around ∩X represent the agent's fallback theories. So, in other words, $P \Rightarrow_X Q$ is true at w iff there is a fallback theory S of X such that S is true at w, S is compatible with P, and S together with P entails Q.

The above definition of $\$_{X,w}$ in terms of $\$_X$ is by no means the only one that makes $\$_{X,w}$ a system of spheres around w satisfying the compatibility condition (nor need it be the most intuitive one). Another such definition would be:

$$\$_{X,w} = \{Y \subseteq W: \text{for some } S \in \$_X, Y = S \cup \{w\}\}.$$

However, from this definition we could prove the rather unintuitive condition:

$$w \in P \Rightarrow_X Q \text{ iff (i) } Q \in X * P \text{ and } w \in -P \cup Q; \text{ or (ii) } \cap X \cap P = \emptyset \text{ and } w \in P \cap Q,$$

saying that the proposition $P \Rightarrow_X Q$ is true at w iff either (i) P materially implies Q at w and Q belongs to the revision of X with P; or (ii) P and Q are both true at w and P is incompatible with the belief state X.

[19] We found this quote in Sahlin (1990). The reference he gives is to *Collected Papers of Charles Sanders Peirce*, ed. Charles Hartshorne and Paul Weiss, Cambridge, Mass. 1931-5, vol. 5, § 27.

[20] Gärdenfors (1988), p. 159f. We have slightly adjusted Gärdenfors' notation in order to make it conform to that in the present paper.

[21] Sahlin (1990), Chap. 4, Section "Conditionals and the Ramsey Test".

[22] Our example is inspired by Gärdenfors' own counterexample to *(RT⇐)* on p. 159 of *Knowledge in Flux* (his Victoria-example).

[23] The conditional in question is "If Hitler had decided to invade England in 1940, Germany would have won the war". Contrast it with the epistemic conditional which I also accept: "If Hitler did decide to invade England in 1940, then he either failed to implement his decision or the invasion somehow misfired".

[24] Our first approach — the one replacing (RT) by Strict (RT) — was constructed just with this idea in mind: to keep *(RT⇒)* while abandoning *(RT⇐)*. The indexical approach gave us both directions, but can be modified so that it only yields *(RT⇐)*. This is accomplished by replacing Compatibility by the following weaker constraint:

(Weak Compatibility)

For every non-empty belief state X, there is a world $w \in \cap X$ such that $\$_{X,w} = \$_X$.

REFERENCES

Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985) 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic* **50**, 510-530.

Chisholm, R. (1946) 'The Contrary-to-fact Conditional', *Mind* **55**, 289-307, reprinted in H. Feigl and W. Sellars (eds.), *Readings in Philosophical Analysis*, Appleton-Century-Crofts, New York 1949, 482-497.

Christensen, D. (1991) 'Clever Bookies and Coherent Beliefs', *Philosophical Review* **100**, 229-248.

Dudman, V. H. (1984) 'Parsing if-sentences', *Analysis* **44**, 145-186.

Dudman, V. H. (1988) 'Indicative and subjunctive', *Analysis* **48**, 114-122.

Dudman, V. H. (1991) 'Interpretations of 'If'-sentences', in Frank Jackson (ed.), *Conditionals,* Oxford University Press, 202-232.

Grove, A. (1988) 'Two modellings for theory change', *Journal of Philosophical Logic* **17**, 157-170.

Gärdenfors, P. (1988) *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press.

Gärdenfors, P., Lindström, S., Morreau, M., and Rabinowicz, W. (1990) 'The negative Ramsey test: another triviality result', in Fuhrmann and Morreau (eds.), *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence no. **465**, Springer-Verlag, 129-134.

Gärdenfors, P. and Makinson, D. (forthcoming) 'Nonmonotonic Inference based on Expectations', to appear in *Artificial Intelligence.*

Hansson, S. O. (1992) 'In Defence of the Ramsey Test', *Journal of Philosophy.* October 1992, 522-540.

Harman, G. (1986) *Change in View: Principles of Reasoning.* Bradford Books, MIT Press.

Harper, W. L. (1976a) 'Rational Belief Change, Popper Functions, and Counterfactuals', in W. L. Harper and C. A. Hooker (eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, vol. 1, Reidel, 73-112.

Harper, W. L. (1976b) 'Ramsey Test Conditionals and Iterated Belief Change (A Response to Stalnaker)', in W. L. Harper and C. A. Hooker (eds.), *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, vol. 1, Reidel, 117.

Jackson, F. (ed.), (1991) *Conditionals*, Oxford Readings in Philosophy, Oxford.

Levi, I. (1988) 'Iteration of conditionals and the Ramsey test', *Synthese* **76**, 49-81.

Lewis, D. (1973) *Counterfactuals*, Basil Blackwell.

Lindström, S. (to appear) 'The Ramsey test and the indexicality of conditionals — a proposed resolution of Gärdenfors' paradox', forthcoming in *Logic, Action and Information,* ed. by A. Fuhrmann and H. Rott, de Gruyter, Berlin-New York.

Lindström, S., and Rabinowicz, W. (1989) 'On probabilistic representation of non-probabilistic belief revision', *Journal of Philosophical Logic* **18**, 69-101.

Lindström, S., and Rabinowicz, W. (1990) 'Epistemic Entrenchment with Incomparabilities and Relational Belief Revision', in Fuhrmann and Morreau (eds.), *The Logic of Theory Change*, Lecture Notes in Artificial Intelligence no. **465**, Springer Verlag, 93-126.

Lindström, S., and Rabinowicz, W. (1992) 'Belief Revision, Epistemic Conditionals and the Ramsey Test', *Synthese* **91**, 195-237.

Quine, W. V. O. (1962) *Methods of Logic*, Second Edition Revised. Routledge & Kegan Paul. London.

Rabinowicz, W. (1991) 'Epistemic Conditionals and Belief Revision', *Annales Academiae Regiae Scientarium Upsaliensis 1989-90,* Almqvist & Wiksell International, Stockholm, 97-112.

Rabinowicz, W. (to appear) 'Stable Revision, or Is Preservation Worth Preserving?', forthcoming in *Logic, Action and Information*, ed. by A. Fuhrmann and H. Rott, de Gruyter, Berlin-New York.

Ramsey, F. P. (1931) *The Foundations of Mathematics and other Logical Essays,* edited by R. B. Braithwaite, Routledge and Kegan Paul, New York.

Ramsey, F. P. (1990) *Philosophical Papers*, edited by D. H. Mellor, Cambridge University Press, Cambridge.

Sahlin, N.-E. (1990) *The Philosophy of F. P. Ramsey*, Cambridge University Press, Cambridge.

Stalnaker, R. (1968) 'A Theory of Conditionals', *Studies in Logical Theory, American Philosophical Quarterly,* Monograph: 2, 98-112. Reprinted in Jackson (1991).

Stalnaker, R. (1975) 'Indicative Conditionals', *Philosophia* 5, 269-86. Reprinted in Jackson (1991).

Teller, P. (1973) 'Conditionalization and Observation', *Synthese* **26**, 218-58.