New Surprises for the Ramsey Test*

Malte Willer

Abstract

In contemporary discussions of the Ramsey Test for conditionals, it is commonly held that (i) supposing the antecedent of a conditional is adopting a potential state of full belief, and (ii) Modus Ponens is a valid rule of inference. I argue on the basis of Thomason Conditionals (such as 'If Sally is deceiving, I do not believe it') and Moore's Paradox that both claims are wrong. I then develop a double-indexed Update Semantics for conditionals that takes these two results into account while doing justice to the key intuitions underlying the Ramsey Test. The semantics is extended to cover some further phenomena, including the recent observation that epistemic modal operators give rise to something very like, but also very unlike, Moore's Paradox.

1 Introduction

How to evaluate a conditional? In 1929, Ramsey made the following suggestion:

> If two people are arguing 'If p will q?' and both are in doubt as to p, they are adding p hypothetically to their stock of knowledge and arguing on that basis about q; so that in a sense 'If p, q' and 'If p, $\overline{q}$' are contradictories. We can say that they are fixing their degrees of belief in q given p.

Ramsey's suggestion concerns open conditionals – those the antecedent of which can be consistently added to what the evaluating agent currently believes. Hence it does not provide a general recipe for evaluating the majority of counterfactual conditionals. These limitations notwithstanding, the quoted passage has served as a starting point for a variety of approaches to the formal semantics of open conditionals. Thus a very popular interpretation runs as follows:

(RT) $A \Rightarrow B \mathbin{@} K$ iff. $B \in K + A$

where K is the set of sentences fully believed true by an agent, and $K + A$ is the result of adding A hypothetically to K. Since it is controversial how the notion of acceptance (denoted here by '@') of a conditional is to be cashed out,

*Published in *Synthese* 176(2): 291–309, 2010

it is left at an intuitive level for the moment.[1] Also, it is inessential to present concerns if and how Ramsey's method can be extended to non-open conditionals. Unless stated otherwise, it will thus be assumed that the antecedent can be consistently added to K.

In critical discussions of the Ramsey Test, so-called 'Thomason Conditionals' like:[2]

(T) If Sally is deceiving me, I do not believe it (because she is so clever).

and the phenomenon of Moore's Paradox, i.e. the unacceptability of constructions of the form:[3]

(M) # ϕ and I do not believe that ϕ.

have served to refute rather crude interpretations of Ramsey's original suggestion. Thus Chalmers. and Hájek (2007) argue that Ramsey's method must not be interpreted as follows:

> 'If ϕ then ψ' is acceptable to an agent S iff., were S to accept ϕ and consider ψ, S would accept ψ.

Due to (M), a rational agent cannot accept that Sally is deceiving him, and then refuse to accept that he believes that Sally is deceiving him. Thus (T), though perfectly rational to assert in some contexts, is predicted to be unacceptable. Since this problem extends to sentences of arbitrary complexity, we get the even worse result that rational agents must take themselves to be omniscient and infallible. Thus the reading is untenable.[4]

According to the account which Chalmers and Hájek – correctly – criticise, an agent who evaluates a conditional should consider the closest possible worlds in which he accepts the antecedent, and check whether these are worlds in which he also accepts the consequent. The reader might have noticed that this is not what Ramsey had in mind, as he suggested that the antecedent is not accepted, but only hypothetically added to what the agent believes to be true. To see the difference, we might give the following reformulation of Ramsey's original suggestion in terms of counterfactuals: To evaluate a conditional, the agent should

[1] The controversy I have in mind is, of course, whether acceptance of $A \Rightarrow B$ at a corpus K is reducible to elementhood of $A \Rightarrow B$ in K. See, e.g., Levi (1977, 1980) and Gärdenfors (1986) for discussion.

[2] The following is a variant of Richmond Thomason's original example as discussed by van Fraassen (1980). What is characteristic of Thomason Conditionals is that the consequent asserts the agent's ignorance or disbelief of the fact described in the antecedent. Two other examples are from Lewis (1986) and Jackson (1987), respectively:

> If Reagan works for the KGB, I'll never believe it.
>
> If Reagan is bald, no one outside his immediate family knows it.

[3] As usual, '#' represents unassertability, which is a much broader than notion than the one of, say, ungrammaticality.

[4] Similar objections to the Ramsey Test thus understood have been raised by Jackson (1987), Edgington (1995), Woods (1997), and Bennett (2003).

consider the closest possible worlds in which he *hypothetically* accepts that the antecedent holds, and check whether he *hypothetically* accepts the consequent in those worlds.[5] Indeed, Arló Costa and Levi (1986) observe that any interesting implementation of the Ramsey Test has to satisfy the following criteria, which are explicitly required by Ramsey himself:[6]

1. The conditionals considered acceptable according to the Ramsey Test are neither truth-value bearers nor objects of belief.

2. The conditionals 'If A, then B', and 'If A, then $\neg B$', cannot be simultaneously acceptable relative to the epistemic state of any agent that is in suspense about A.

3. The conditionals delivered by the Ramsey Test are to be understood as expressions of suppositional reasoning.

4. An agent who is in suspense about A accepts 'If A, then B' with respect to his epistemic state iff B belongs to the belief state obtained after adding A to that state.

Since the third criterion clearly requires that the antecedent is only supposed to be true by the evaluator of a conditional, neither Thomason Conditionals nor Moore's Paradox seem to be of high interest for the status of the Ramsey Test.

This view, however, underestimates the importance of Thomason Conditionals and Moore's Paradox for the Ramsey Test. Taken together, they have much greater philosophical momentum than what has so far been realised. In particular, they show that (i) supposing cannot, as traditionally assumed, be reconstructed in terms of adopting a potential belief state and (ii) Modus Ponens is an invalid rule of inference. These issues will be discussed in turn. As

[5]Assuming Strong Centering, this method of evaluating conditionals is fully equivalent to the suggestion made in the quote from Ramsey, taken literally: Add the antecedent hypothetically to what you believe, and see whether you then hypothetically accept the consequent. I see no benefit in a counterfactual phrasing of Ramsey's method, and thus will stick to a literal reading of his suggestion in what follows.

[6]That Ramsey required the first criterion is indicated by various passages in his 'General Propositions and Causality'. Conditionals – which Ramsey also calls 'variable hypotheticals' – are not believed, but encode rules for belief formation (page reference is to the reprint in Ramsey (1990)):

> Variable hypotheticals are not judgements but rules for judging 'If I meet a ϕ, I shall regard it as a ψ'. This cannot be *negated* but it can be *disagreed* with by one who does not adopt it. (p. 149)

Ramsey held that conditionals are not truth-value bearers since we can agree on all the facts and still disagree about a conditional. He considers a scenario in which we disagree with someone who holds that if he eats the cake, he will get sick, and thus does not eat it:

> After the event we both know that he did not eat the cake and that he was not ill; the difference between us is that he thinks that if he had eaten it he would have been ill, whereas we think he would not. But this is *prima facie* not a difference of degrees of belief in any proposition, for we both agree as to all the facts. (p. 155)

That Ramsey required the other criteria is clear from the quote in §1.

a final contribution, this paper presents a formal framework for the semantics of conditionals which takes these lessons into account and thus constitutes a substantial improvement over former implementations of the Ramsey Test.

2 Supposing and Potential States of Belief

How to make a supposition? The intuitive answer is that in supposing A, an agent temporally adopts the belief state he would be in, were he to accept that A is true. This answer is not only intuitive and widely accepted, but also plays an important role in explaining the important role of suppositions for practical and theoretical deliberations. Thus Levi observes that in practical deliberation, rational agents can, and in fact often do, suppose that a certain option is chosen without accepting, or coming to believe, that the option is implemented. In this context, he writes:

> The state of full belief reached by adding the information that the option is chosen and is going to be implemented is the same potential state of full belief as the suppositional state reached by supposing for the sake of the argument that the option is implemented. That is why suppositional reasoning is so useful in practical deliberation. Levi (1996, p. 4)

Levi also points out that rational inquiry regularly involves supposing a certain answer and exploring its consequences *before* the answer is believed to be true. Furthermore, changing a state of belief requires some sort of accounting of justification, while supposition does not. Critical assessment of the latter is restricted to the question to what degree, if any, the supposition is relevant for the problem at hand. But also in the context of theoretical deliberation, Levi claims that supposing is adopting a potential belief state (where $\underline{\mathrm{K}}$ is the agent's current belief state and h is a sentence):

> If the inquirer subsequently adopts the potential answer represented by expanding $\underline{\mathrm{K}}$ through adding h as his or her full state of full belief, the formal structure of the shift from $\underline{\mathrm{K}}$ is the same as when $\underline{\mathrm{K}}$ is expanded by supposing h for the sake of the argument. (Ibid.)

Levi gives voice to a strongly motivated view on suppositional reasoning: It is the match between the potential state of full belief that A and the epistemic state reached by supposing that A which explains for Levi the usefulness of suppositional reasoning. If the two were different – if supposing A would result in an epistemic state notably different from the epistemic situation the agent would be in were he to genuinely believe A – how could suppositional reasoning be useful for rational inquiry? After all, suppositional reasoning is designed to show us the vices and virtues of adopting a certain potential state of full belief – unless suppositional reasoning is a reliable guide to states of full belief, its usefulness for rational inquiry appears unexplained.

These are philosophical considerations which need to be taken seriously. But there is an argument which strongly points in the opposite direction. The view that supposing is adopting a potential belief state is inconsistent with some plausible assumptions about belief states and hypothetical reasoning. To develop the point in detail, epistemic states will be modelled, as usual, as sets of sentences closed under the logical consequence relation *Cn*, i.e. theories. The epistemic state of an agent comprises everything he is committed to qua his avowed beliefs and qua logical consequence – regardless of whether the agent does or even can live up to his commitments. The language is confined to classical propositional language, but extended with constructions involving the Ramsey conditional ($\Rightarrow$) and belief self-ascriptions of the form 'I believe that ϕ' ($\textsc{Bel}(\phi)$). Define:

> $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$ is the smallest set containing any sentential atoms $\mathcal{A} = \{p, q, ...\}$ and is closed under negation ($\neg$), conjunction ($\wedge$), the Ramsey conditional ($\Rightarrow$) and the belief operator ($\textsc{Bel}$).

The following minimal constraints on how a rational agent ought to move from one epistemic state to another by supposing that ϕ apply:

> (Success) For all $\phi \in \mathcal{L}^{\mathsf{B}}_{\Rightarrow}$: $\phi \in K + \phi$
>
> (Consistency) If $\neg\phi \notin Cn(\varnothing)$, then $K + \phi$ is consistent

Success requires that posterior states carry commitments to the information inducing the change; Consistency insists that hypothetical reasoning should be consistency preserving and, where this does not conflict with Success, consistency restoring. These constraints should be familiar from discussions of belief revision models, and for present purposes I treat them as non-negotiable.

A belief set is a theory that reflects our interests in agents who believe the (classical) consequences of what they believe. Not just any theory will do as a belief set. For present purposes, we want agents to be reflective, their belief sets encoding not only first-order beliefs but also being closed under what the agent considers to be his own doxastic state. Adopting a standard procedure, we capture this by first defining an operation Dox as follows:

> If K is a corpus, then the meta-corpus $\text{Dox}(K)$ is the smallest set such that:
>
> 1. if $\phi \in K$, then $\textsc{Bel}(\phi) \in \text{Dox}(K)$
> 2. if $\phi \notin K$, then $\neg\textsc{Bel}(\phi) \in \text{Dox}(K)$
>
> A belief set K is *closed under* Dox iff $\text{Dox}(K) \subseteq K$.

Reflectivity can now be put as the constraint that we only consider Dox-closed belief sets:

> (Reflectivity) If K is a belief state, then K is closed under Dox.

Closure of K under Dox is not only intuitive when it comes to belief, but also supported by a doxastic interpretation of Moore's Paradox. In order to accept a conjunction, an agent must accept both conjuncts. Hence a rational agent who accepts (M) accepts that ϕ and that he does not believe that ϕ. Thus both conjuncts are in the agent's belief state. But the belief that ϕ is true and Reflectivity commit the agent to believing that he believes that ϕ. Reflectivity thus straightforwardly explains why (M) is unacceptable to a rational agent: accepting (M) would commit the agent to holding incompatible beliefs about his doxastic situation.[7]

Our epistemic capacities are limited. There are truths we do not believe. We assume that some rational agents are *modest* in the sense that they are aware of these limitations:[8]

(Modesty) For some K, for some $\neg A \notin Cn(\varnothing) : A \Rightarrow \neg \text{Bel}(A) \mathbin{@} K$

Modesty is entailed by the observation that Thomason Conditionals are sometimes perfectly acceptable to rational agents.

It can now be proven that the preceding principles are incompatible with the view that suppositional reasoning is to be identified with adopting a potential belief state, i.e. with the following view about suppositional reasoning:

(SR) For all $\phi \in \mathcal{L}^{\mathsf{B}}_{\Rightarrow} : K + \phi$ is a belief state.

Proof. By Modesty, there is a K such that $A \Rightarrow \neg\text{Bel}(A) \mathbin{@} K$, $\neg A \notin Cn(\varnothing)$. Thus by RT, $\neg\text{Bel}(A) \in K + A$. But by SR, $K + A$ is a belief state, and so $\text{Dox}(K + A) \subseteq K + A$ by Reflectivity. Since $A \in K + A$ by Success, $\text{Bel}(A) \in K + A$. Thus $K + A$ is inconsistent, whence by Consistency $\neg A \in Cn(\varnothing)$. Contradiction.

The view that supposing is adopting a potential state of belief is widely accepted and enjoys some important philosophical motivation. This notwithstanding, the competing principles have stronger support. We do not want to give up Success and Consistency. Reflectivity and Modesty are highly intuitive and, what is more, Reflectivity explains why Moorean constructions are unacceptable, while Modesty is entailed by the fact that Thomason Conditionals are sometimes perfectly acceptable to rational agents. We thus have to re-think what it means to suppose that something is the case. Let me conclude with some final remarks before we move on to Modus Ponens.

There is an interesting resemblance between the present result and what has become known as the 'Fuhrmann Impossibility Result'.[9] To illustrate this point,

[7] Shoemaker (1995) suggests that what can be coherently asserted is constrained by what can be coherently believed. If this is correct, it follows as a corollary that it is absurd to assert (M).

[8] As before, we should avoid equating acceptance of $A \Rightarrow B$ at K with elementhood of $A \Rightarrow B$ in K. Hence the use of '@' instead of '$\in$'.

[9] See Levi (1988), Fuhrmann (1989), Rott (2001). The term 'Fuhrmann Impossibility Result' goes back to Hansson (1999).

extend $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$ with constructions involving the epistemic modals $\Diamond$, $\Box$, understood as duals. Thus we add to $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$ constructions stating what *might* and *must* be the case. Analogous to what happened before, we first define an operation Poss as follows:

> If K is a corpus, then the meta-corpus Poss(K) is the smallest set such that:
>
> 1. if $\phi \in K$, then $\Box\phi \in \text{Poss}(K)$
> 2. if $\neg\phi \notin K$, then $\Diamond\phi \in \text{Poss}(K)$
>
> A belief set K is *closed under* Poss iff $\text{Poss}(K) \subseteq K$.

The Fuhrmann Impossibility Result is originally concerned with belief revision, but straightforwardly applies to hypothetical reasoning. It shows that if '+' obeys Success and Consistency, the following principles of Preservation and Reflectivity* are inconsistent if there are non-trivial corpora K.

> (Preservation) If $\neg\phi \notin K$, then $K \subseteq K + \phi$.
>
> (Reflectivity*) If K is a belief state, then K is closed under Poss.
>
> (Non-Triviality) There is a corpus K such that: neither $\phi \in K$ nor $\neg\phi \in K$

Proof. Given non-trivial K, select A such that neither $A \in K$ nor $\neg A \in K$. Due to Reflectivity, Poss(K)$\subseteq K$ and thus $\Diamond A \in K$, $\Diamond\neg A \in K$. Consider $K+A$. Since $A \notin K$, $K \subseteq K+A$ by Preservation. Thus $\Diamond\neg A \in K+A$. But by Success, $A \in K+A$, and thus due to Closure under Poss, $\Box A \in K+A$. Thus $K+A$ is inconsistent, whence by Consistency $\neg A \in Cn(\varnothing)$. But all belief sets are closed under Cn, so $\neg A \in K$. Contradiction.

In words, this result shows that agents who have beliefs about what might and might not be the case given their current epistemic situation cannot be doxastically preservative in hypothetical reasoning, unless they are fully opinionated about objective matters of fact (the trivial case).

The results are similar since both involve, in addition to Success and Consistency, an introspective principle (Reflectivity, Reflectivity*) and a fairly uncontentious existence claim about corpora (Modesty, Non-Triviality). The present result, however, does not appeal to the controversial Preservation principle, which is rejected by most authors as a reaction to the Fuhrmann Impossibility Result. Instead, the inconsistency of $K+A$ is derived via the Ramsey Test and the assumption that supposing results in a state of belief. In fact, the use of (RT) brings the present observation closer to Gärdenfors's (1986) result that the Ramsey Test and the Preservation Principle are, on pain of triviality, inconsistent with each other. Notice, however, that Gärdenfors equates acceptability of a conditional at a corpus with set membership, i.e. his proof employs (O):

(O) $A \Rightarrow B \in K$ iff. $B \in K + A$

No step in the argument against the claim that supposing is adopting a potential belief state requires (O) – the notion of accepting a conditional is left at an intuitive level.

However, the fact that both results involve some introspective principle might be grist for the mills of those who are sceptical of Reflectivity. Levi denies that we ever have beliefs of serious possibility, banning epistemic modals and conditionals from the corpus K. It seems to me abundantly clear that there are second-order beliefs of the form 'I believe that ϕ'. Such judgements qualify as 'biographical remarks' in Levi's sense – judgements an agent makes about his own state of belief, evaluable for truth or falsity. Thus Levi himself would have to admit that the belief self-ascriptions which figure in the presented result are regular beliefs, admissible into the corpus K. Reflectivity, then, appears to be beyond dispute. Accordingly, we need to give up the assumption that engaging in hypothetical reasoning is adopting a potential belief state.

3 Modus Ponens

What inferences are licensed by conditionals? According to the rule of Modus Ponens, an indicative conditional of the form $\ulcorner \phi \Rightarrow \psi \urcorner$ licenses, together with the antecedent ϕ, inference to the conclusion that ψ. In other words, $\ulcorner \phi \Rightarrow \psi \urcorner$ is at least as strong as the material conditional $\ulcorner \phi \supset \psi \urcorner$, i.e.:

(MP) $\quad \ulcorner \phi \Rightarrow \psi \urcorner \vdash \ulcorner \phi \supset \psi \urcorner$

I am concerned here with Modus Ponens as a rule of inference, not as a law of semantics about the preservation of truth.[10] McGee (1985) has famously offered alleged counterexamples to Modus Ponens thus understood. However, his counterexamples are limited to cases in which the consequent is itself a conditional, making his attack on Modus Ponens susceptible to a series of critical responses.[11] The following counterexample to Modus Ponens is immune to such criticisms and thus constitutes, I think, a substantial improvement over the scepticism about this rule as originally expressed by McGee.

To establish the counterexample to Modus Ponens as a rule of inference, we show that there is an occasion on which one has good grounds for believing the premises of an application of Modus Ponens but yet one is not justified in accepting the conclusion. Specifically, in light of Moore's Paradox, Modus Ponens turns out as inferentially unreliable for modest rational agents. Let me explain.

[10] Remember in this context Arló Costa's and Levi's observation that Ramsey's conditionals are not truth-value bearers at all. As a semantic rule, Modus Ponens is not applicable to conditionals thus understood.

[11] See, e.g., the defenses of Modus Ponens offered by Sinnott-Armstrong et al. (1986) and Lowe (1987)

Moorean paradoxical constructions are unacceptable to agents who are rational. Unacceptability comes in different flavours. The agent might be unable to accept ϕ since he lacks sufficient evidence in support of ϕ. The agent might also be unable to accept ϕ since ϕ is a priori absurd. For instance, ϕ might be an obvious contradiction. In such cases, not only is ϕ unacceptable, but also is the agent rationally committed to accept $\ulcorner \neg\phi \urcorner$. Neither accepting ϕ nor $\ulcorner \neg\phi \urcorner$ is not an option: Unacceptability of ϕ commits any rational agent to acceptance of $\ulcorner \neg\phi \urcorner$. Moorean paradoxical constructions are unacceptable in the latter sense: It is simply absurd for a rational agent to judge true both that ϕ holds and that he does not believe that ϕ holds. Hence rational agents ought to accept the negation of (M):

(N) Not: ϕ and I do not believe that ϕ.

Consider an agent S who accepts the Thomason Conditional 'If Sally is deceiving me, I do not believe that she is deceiving me'. Accepting (MP) yields the following as a valid argument for S, with 'A' abbreviating 'Sally is deceiving me':

(1)	$A \Rightarrow \neg \mathrm{Bel}(A)$	(assumption)
(2)	$\neg(A \wedge \neg \mathrm{Bel}(A))$	(N)
(3)	$A \supset \neg \mathrm{Bel}(A)$	(1, MP)
(4)	$A \supset \mathrm{Bel}(A)$	(2, tautology)
(5)	$\neg A$	(3, 4, tautology)

But this is the wrong result. Certainly, clever women are not always loyal, so S should not be allowed to infer Sally's loyalty from her cleverness. S has very good reason to believe (1) and (2) and to trust the tautological laws involved, but (MP) leads to a conclusion S does not have very good reason to believe.

To avoid potential confusions, notice that the counterexample does not involve a case in which an agent accepts a conditional and rejects it once the antecedent is learnt. Such scenarios are perfectly compatible with Modus Ponens, but irrelevant for present purposes. The problematic inference occurs without the agent learning that Sally is deceiving him, and thus there is no reason for the agent to give up the conditional.

It should be noted that this argument does not make any specific assumptions about the Ramsey Test, but works against any theory of conditionals that predicts (MP) as a valid rule of inference. However, let K/A be the result of coming to believe that A in state K. Then it is a minimal condition for B being in K/A that $A \supset B \in K$ – otherwise K/A could not capture the logical commitments of an agent who comes to believe that A in K. Now remember the assumption that supposing that A is adopting the state one would be in, were one come to believe that A, and put this more rigidly as $K + A = K/A$. Then (MP) follows straightforwardly. In other words, (MP) follows from (RT) and the following principles:

(Commitment) If $\psi \in K/\phi$, then $\phi \supset \psi \in K$

(SR*) For all $\phi \in \mathcal{L}^{\mathsf{B}}_{\Rightarrow}$: $K + \phi = K/\phi$

Proof. Assume that $A \Rightarrow B \,@\, K$. By (RT), $B \in K + A$ and thus $B \in K/A$ by (SR*). Whence $A \supset B \in K$ by Commitment.

If it is agreed that Commitment looks pretty good, then once again it is problematic to assume that supposing A is adopting the potential state of full belief that A. If this assumption is dropped, supposing that A might result in an epistemic state which is different in kind from the one the agent would be in, were he come to believe A. Specifically, and set theoretically speaking, the former need not be a subset of the latter, thus blocking the derivation of Modus Ponens from the Ramsey Test. Once again, we have reason not to reconstruct supposing as adopting a potential belief state.

4 Double-Indexed Dynamic Semantics

How to accommodate the lesson that supposing is not to be understood as adopting a potential belief state? Here is a suggestion: In an abstract representation of information states, one needs to keep track of what the agent believes and what, in addition, is supposed. That is, information states are to be represented as double-indexed, where one index keeps track of what is believed, while the other index keeps track of additional assumptions. In representing information processing, we can now distinguish between supposing and coming to believe. These ideas are most perspicuously elaborated by revising an Update Semantics as originally developed by Veltman (1996).[12]

The Update Semantics is designed for our language $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$, which, to repeat, is the smallest set that contains any sentential atoms $\mathcal{A} = \{p, q, ...\}$ and is closed under negation ($\neg$), conjunction ($\wedge$), the Ramsey conditional ($\Rightarrow$) and the belief operator (Bel).

Definition 1 (Possibilities and Double-Indexed Information States) Fix a set $\mathcal{A}$ of atomic formulas, w is a possibility iff w: $\mathcal{A} \rightarrow \{0, 1\}$. W is the set of such w's. An ordered pair $\langle s_1, s_2 \rangle$ is a double-indexed information state iff $s_1, s_2 \subseteq W$. I is the set of such $\langle s_1, s_2 \rangle$'s. An information state $\langle s_1, s_2 \rangle$ is absurd iff. s_1 or s_2 are the empty set. An information state $\langle s_1, s_2 \rangle$ is regular iff. $s_1 = s_2$.

The notion of an information state as laid down in the first definition differs from the one of the preceding sections in two respects: Sets of sentences are

[12]For further elaborations of Update Semantics, see Groenendijk et al. (1996), van der Does et al. (1997), and Beaver (2001). Gillies uses Update Semantics in his studies of Moore's Paradox (2001), indicative conditionals (2004a), and belief revision (2004b; 2006). The considerations presented in the sections before apply to Gillies's work without loss of generality.

replaced with sets of possible worlds, and some structure is added via double-indexing. The first difference is merely notational; the second, substantial one, allows keeping track of what is believed and assumed. I stipulate that s_2 keeps track of what is believed, while s_1 keeps track of what the agent assumes. If $s_1 = s_2$, the information that is relevant for the agent's hypothetical reasoning is fully provided by what he believes. If an agent with information state $\langle s_1, s_2 \rangle$ assumes that p holds yet does not believe it, p will be true in all possibilities of s_1 but not of s_2. For current purposes, we are interested in those cases in which $s_1 \subseteq s_2$, i.e. in those cases in which the agent's assumptions are not in conflict with his beliefs. It is worth noticing, however, that there is no formal obstacle to allowing for free variation of beliefs and assumptions.

In information processing, we update our information state, i.e. updating is a process which takes you from one information state to the next given some new input. An agent can hypothetically update her information state with ϕ without assuming that she believes that ϕ. Double-indexing allows us to capture this point: Hypothetically accepted sentences induce updates on the first position in the agent's information state, but what matters for the correctness of a belief ascription is the second position of that information state.

Definition 2 (Hypothetical Updates on Double-Indexed Information States) Consider any $w \in W$, $\langle s_1, s_2 \rangle, \langle s_3, s_4 \rangle \in I$, $p \in \mathcal{A}$, and $\phi, \psi \in \mathcal{L}^{\mathsf{B}}_{\Rightarrow}$. Define a function $\boxminus : I \times I \to I$ in the following way: $\langle s_1, s_2 \rangle \boxminus \langle s_3, s_4 \rangle = \langle (s_1 \backslash s_3), s_2 \rangle$. A hypothetical update $\uparrow$ is a function: $I \to I$ defined by the following recursion:

(1) $\langle s_1, s_2 \rangle \uparrow p = \langle \{w \in s_1 : w(p) = 1\}, s_2 \rangle$;
(2) $\langle s_1, s_2 \rangle \uparrow \neg\phi = \langle s_1, s_2 \rangle \boxminus (\langle s_1, s_2 \rangle \uparrow \phi)$;
(3) $\langle s_1, s_2 \rangle \uparrow (\phi \wedge \psi) = (\langle s_1, s_2 \rangle \uparrow \phi) \uparrow \psi$;
(4) $\langle s_1, s_2 \rangle \uparrow (\phi \Rightarrow \psi) = \langle \{w \in s_1 : (\langle s_1, s_2 \rangle \uparrow \phi) \uparrow \psi = \langle s_1, s_2 \rangle \uparrow \phi\}, s_2 \rangle$;
(5) $\langle s_1, s_2 \rangle \uparrow \text{Bel}\ (\phi) = \langle \{w \in s_1 : \langle s_2, s_2 \rangle \uparrow \phi = \langle s_2, s_2 \rangle\}, s_2 \rangle$.

A rational agent should at least hypothetically accept ϕ if he is in an information state $\langle s_1, s_2 \rangle$ such that $\langle s_1, s_2 \rangle \uparrow \phi = \langle s_1, s_2 \rangle$.

Remember the stipulation that s_2 keeps track of what the agent believes. To see what is going on in the second definition, notice that a hypothetical update on an information never induces a change on s_2. This captures the fact that assuming that ϕ is the case does not affect what the agent believes. But as one can see from (1), assuming that Sally is deceiving me is *quite like* coming to believe that Sally is deceiving me, in the sense that it removes epistemic uncertainty: (1) requires that hypothetically updating an information state with an atom p eliminates all possibilities from s_1 in which p is false. Clause (2) integrates Veltman's original treatment of negation as set subtraction into double-indexed dynamic semantics: In hypothetically updating an information state $\langle s_1, s_2 \rangle$ with $\ulcorner \neg\phi \urcorner$, we first determine the result of hypothetically updating that state with ϕ, $\langle s_1, s_2 \rangle \uparrow \phi$, and then subtract the first position of the latter from the first position of the original state (this is just what '$\boxminus$' does). So for example, hypothetical update with $\neg p$ only leaves those possibilities in s_1 in which p is false, but s_2 unmodified. To hypothetically update an information state with a conjunction, an agent first hypothetically updates with the first conjunct, and then hypothetically updates the resulting state with the second conjunct: conjunction is functional composition (cf. (3)).

Clause (4) encodes Ramsey's suggestion for evaluating conditionals. A conditional invites an agent to perform a *test* on his information state. If hypothetically updating the information state with ϕ and then with ψ would add no more information than a mere hypothetical update with ϕ, i.e. if $(\langle s_1, s_2\rangle \uparrow \phi) \uparrow \psi = \langle s_1, s_2\rangle \uparrow \phi$, the state passes the test posed by the conditional $\ulcorner\phi \Rightarrow \psi\urcorner$. Put another way, an agent's information state passes the test posed by $\ulcorner\phi \Rightarrow \psi\urcorner$ if and only if the agent will hypothetically accept ψ under the assumption that ϕ. Clause (4) then simply combines this idea with our previous agreement that hypothetical updates only affect the first position of an agent's information state: Hypothetical update of $\langle s_1, s_2\rangle$ with a formula of the form $\ulcorner\phi \Rightarrow \psi\urcorner$ returns the original state $\langle s_1, s_2\rangle$ if it holds that $(\langle s_1, s_2\rangle \uparrow \phi) \uparrow \psi = \langle s_1, s_2\rangle \uparrow \phi$. Otherwise, the output is $\langle \varnothing, s_2\rangle$.

Now let's move on to clause (5). Quite like a Ramsey conditional, a belief self-ascription is concerned with the information state of the evaluating agent. Self-ascribing the belief that ϕ is adequate just in case the ascriber already believes the information encoded in ϕ. In our framework, that will hold just in case the information that ϕ carries is already contained in the information encoded in s_2: Adding the information carried by ϕ to s_2 produces no change at all to s_2. This intuition is clear enough, but we need to overcome a small technical difficulty: To check whether the agent believes that ϕ, it would *not* be sufficient to check whether $\langle s_1, s_2\rangle \uparrow \phi = \langle s_1, s_2\rangle$. This condition is fulfilled just in case the information carried by ϕ is already contained in the information encoded in s_1 – we are simply looking at the wrong index. Of course, no problem occurs in case $s_1 = s_2$, but the main point of the framework under investigation is that an agent can add information to s_1 without adding it to s_2. Fortunately, a simple modification will do the trick: We first copy s_2 into the first position, and then check whether $\langle s_2, s_2\rangle \uparrow \phi = \langle s_2, s_2\rangle$. That condition is fulfilled just in case the information carried by ϕ is already contained in the information encoded in s_2 – we are guaranteed to look at the right index. If we combine this idea with the general treatment of hypothetical updates, this gives us clause (5): Hypothetical update of $\langle s_1, s_2\rangle$ with a formula of the form $\ulcorner\text{Bel}\ (\phi)\urcorner$ returns the original state $\langle s_1, s_2\rangle$ just in case $\langle s_2, s_2\rangle \uparrow \phi = \langle s_2, s_2\rangle$. Otherwise, the output is $\langle \varnothing, s_2\rangle$.[13] I will now describe what happens when an agent *accepts* a formula ϕ, i.e. when he *genuinely updates* his information state with ϕ.

[13]The current treatment of hypothetical updates as tests as laid out in (4) and (5) is good enough for present purposes. In §5.3, I will say more about what might happen in case an agent's information state fails such a test.

In order to prepare the introduction of the notion of accepting (as opposed to supposing), I introduce a simple function which brings the second position of an information state in line with the first position of that state.

Definition 3 (The (*)-Function) Define a function $^*: I \to I$:

(*) $\langle s_1, s_2 \rangle^* = \langle s_1, s_1 \rangle$.

We are now in the position to capture the notion of acceptance by defining the notion of a genuine update (as opposed to a mere hypothetical update) on an information state. The idea is very simple: Genuine update with ϕ is quite like hypothetical update with ϕ, the difference being that the former also affects what the agent believes, i.e. the second position of his information state. This is where the (*)-function gets important: It guarantees that the effects of hypothetically accepting ϕ on s_1 get carried across to the second position.

Definition 4 (Genuine Updates on Double-Indexed Information States) Consider any $\langle s_1, s_2 \rangle \in I$, $p \in \mathcal{A}$, and $\phi, \psi \in \mathcal{L}^{\mathsf{B}}_{\Rightarrow}$. A genuine update function $\uparrow^*: I \to I$ is defined by the following recursion:

(i) $\langle s_1, s_2 \rangle \uparrow^* p = (\langle s_1, s_2 \rangle \uparrow p)^*$;
(ii) $\langle s_1, s_2 \rangle \uparrow^* \neg\phi = (\langle s_1, s_2 \rangle \boxminus (\langle s_1, s_2 \rangle \uparrow^* \phi))^*$;
(iii) $\langle s_1, s_2 \rangle \uparrow^* (\phi \wedge \psi) = (\langle s_1, s_2 \rangle \uparrow^* \phi) \uparrow^* \psi$;
(iv) $\langle s_1, s_2 \rangle \uparrow^* (\phi \Rightarrow \psi) = (\langle s_1, s_2 \rangle \uparrow (\phi \Rightarrow \psi))^*$;
(v) $\langle s_1, s_2 \rangle \uparrow^* \text{BEL}\ (\phi) = (\langle s_1, s_2 \rangle \uparrow \text{BEL}\ (\phi))^*$.

To get a feel for the mechanism, consider (i): An agent who accepts p first hypothetically updates with p, thus eliminating the possibility of p being false in s_1. We then copy s_1 over into the second position, thus guaranteeing that the possibility of p being false is eliminated in both indices. The resulting information state now passes the test posed by ⌜BEL (p)⌝, which is exactly the result we wanted. It is easy to verify that the same mechanism is at work in the other clauses.

The final step is to introduce the notions of support, entailment and equivalence. Intuitively, an information state commits you to ϕ just in case the information carried by ϕ is already contained in the information encoded in that information state. One way of making this more precise is this: An information state $\langle s_1, s_2 \rangle$ commits you to ϕ, ϕ is supported by $\langle s_1, s_2 \rangle$, just in case genuine update of that state with ϕ amounts to nothing more than a genuine update of that state with any tautology.

Definition 5 (Support, Entailment, Equivalence) Let $\langle s_1, s_2\rangle$ be an information state, and ϕ, ψ be any formulas in $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$. The support and commitment relation $\vDash$ is defined as follows (where $\mathbf{1}$ is any tautology):[14]

(I) $\langle s_1, s_2\rangle$ *supports* ϕ, $\langle s_1, s_2\rangle \vDash \phi$, iff $\langle s_1, s_2\rangle \uparrow^* \phi = \langle s_1, s_2\rangle \uparrow^* \mathbf{1}$;

(II) ϕ *commits to* ψ (ϕ entails ψ), $\phi \vDash \psi$, iff it holds for all $\langle s_1, s_2\rangle$ that: $\langle s_1, s_2\rangle \uparrow^* \phi \vDash \psi$;

(III) ϕ, ψ are *epistemically equivalent*, $\phi \cong \psi$, iff it holds for all $\langle s_1, s_2\rangle$ that: $\langle s_1, s_2\rangle \uparrow^* \phi = \langle s_1, s_2\rangle \uparrow^* \psi$.

Let me make explicit why the current revision model $\langle I, \uparrow, \uparrow^*, \vDash\rangle$ is a substantial improvement over previous approaches to the Ramsey Test.

First, supposing is not reconstructed in terms of adopting a potential belief state. The epistemic state an agent is in when he supposes ϕ may be crucially different from the one he would be in, were he to accept ϕ. In the former case, only the first position of the epistemic state is modified; in the latter, both positions may undergo a change.

Second, it becomes immediately clear from (1) and (5) that hypothetically updating an information state with a sentence ϕ does not guarantee that the information state passes the test posed by $\ulcorner\text{Bel}\ (\phi)\urcorner$. In other words, an agent can coherently assume $\ulcorner\phi \wedge \neg\text{Bel}\ (\phi)\urcorner$. And this is all we need to predict that a Thomason Conditional of the form $\ulcorner\phi \Rightarrow \neg\text{Bel}\ (\phi)\urcorner$ may be perfectly acceptable to a rational agent. Remember what we said about the Ramsey conditional: An agent's information state passes the test posed by such a conditional just in case the agent hypothetically accepts the consequent on the assumption that the antecedent is true. So an agent accepts a Thomason Conditional just in case he hypothetically accepts $\ulcorner\neg\text{Bel}\ (\phi)\urcorner$ on the assumption that ϕ. And this is a real possibility given the current framework. To offer a simple example, take an agent's regular information state $\langle s_1, s_2\rangle$, so that both p and $\neg p$ are possibilities in s_1, s_2. As it stands, $(\langle s_1, s_2\rangle \uparrow p) \uparrow \neg\text{Bel}\ (p) = \langle s_1, s_2\rangle \uparrow p$, and thus $\ulcorner p \Rightarrow \neg\text{Bel}\ (p)\urcorner$ is acceptable to the agent. Coming back to the case of Sally, (T) is acceptable by a rational agent in those cases in which the agent does not already believe that Sally is deceiving him and has no good reason to believe that he would become informed about Sally's future attempts of deceiving him.

Third, Moore's Paradox is nevertheless fully captured by the present framework. For each $\langle s_1, s_2\rangle$, $\langle s_1, s_2\rangle \uparrow^* \phi \vDash \text{Bel}\ (\phi)$ and thus $\langle s_1, s_2\rangle \vDash \neg(\phi \wedge \neg\text{Bel}\ (\phi))$. So even though $\ulcorner\phi \wedge \neg\text{Bel}\ (\phi)\urcorner$ can be coherently *assumed*, it cannot be coherently *accepted*. Another way of putting the point is that the result

[14] To see why updating with $\mathbf{1}$ is advantageous in (I), consider the major alternative:

$\langle s_1, s_2\rangle$ *supports* ϕ, $\langle s_1, s_2\rangle \vDash \phi$, iff, $\langle s_1, s_2\rangle \uparrow^* \phi = \langle s_1, s_2\rangle$.

According to this definition, an irregular information state (i.e. if $s_1 \neq s_2$) does not support anything, not even a tautology. The (*)-function brings the first and second position in line, and thus induces a change on all irregular information states. The definition in (I) takes care of this problem.

of accepting ϕ in state $\langle s_1, s_2 \rangle$ is reflective in the sense that $\phi \vDash$ Bel (ϕ). In contrast, the result of hypothetically accepting ϕ is not always reflective since there are $\langle s_1, s_2 \rangle$ such that $\langle s_1, s_2 \rangle \uparrow \phi \nvDash$ Bel (ϕ). It is at this point that the inconsistency of §2 is avoided: If ϕ is accepted in a state $\langle s_1, s_2 \rangle$, i.e. if $\langle s_1, s_2 \rangle \vDash \phi$, that state is reflective in the sense that $\langle s_1, s_2 \rangle \vDash$ Bel (ϕ). But the state reached by hypothetically accepting ϕ (corresponding to $K + \phi$ from §§ 2, 3) is not always a state in which ϕ is accepted. Specifically, there are $\langle s_1, s_2 \rangle$ such that $\langle s_1, s_2 \rangle \uparrow \phi \nvDash \phi$ and $\langle s_1, s_2 \rangle \uparrow \phi \nvDash$ Bel (ϕ).

Fourth, the current framework predicts that rational commitment is not closed under the rule of Modus Ponens (MP), in the sense that an agent may be committed to $\ulcorner \phi \Rightarrow \psi \urcorner$ without being committed to $\ulcorner \phi \supset \psi \urcorner$. This follows immediately from the previous remarks. The conclusion of the second remark was that there are $\langle s_1, s_2 \rangle$ such that $\langle s_1, s_2 \rangle \vDash \phi \Rightarrow \neg$Bel (ϕ). The conclusion of the third remark was that for each $\langle s_1, s_2 \rangle$, $\langle s_1, s_2 \rangle \vDash \neg(\phi \wedge \neg$Bel $(\phi))$. Thus for non-trivial ϕ, $\phi \Rightarrow \neg$Bel $(\phi) \nvDash \phi \supset \neg$Bel (ϕ). Again, this should come as no surprise, given the distinction between supposing that ϕ and accepting ϕ.

Fifth, it was observed in §2 that we need to explain the usefulness of suppositional reasoning for practical and theoretical deliberation. Such an explanation can be provided even without identifying suppositional reasoning with adopting a potential belief state. An agent who supposed ϕ and later accepts ϕ comes to believe what he originally merely assumed. But this is not to say that he actually adopts the epistemic state he was in when he engaged in hypothetical reasoning. Rather, the information which originally played a role only for his hypothetical reasoning now becomes available for operations which are determined by what the agent believes. In the present framework, this is represented by the (*)-operation: The information encoded in s_1 as strengthened by ϕ gets 'copied across', thus being now included in both s_1 and s_2. Hypothetical reasoning can be understood as a valuable tool for practical and theoretical deliberation even though supposing that ϕ results in an epistemic state which is different from the one resulting from coming to believe that ϕ.

5 Further Elaborations

5.1 In the Neighbourhood of Modus Ponens

It was observed that Modus Ponens is an invalid rule of inference. However, there is the strong intuition that something in the neighbourhood of it must hold in a logic of rational belief change. As Gillies (2004a) puts it, a rational agent is committed to $\ulcorner \phi \Rightarrow \psi \urcorner$ just in case learning that ϕ would commit him directly to ψ. Given that 'S learns ϕ' means that ϕ is true and S is informed about that, I suggest a slight modification of Gillies's intuition: Assume that learning ϕ would commit me directly to ψ. I then should accept that 'If I learn that ϕ, then (I accept) ψ'. To capture this conditional construction, I add another binary connective ($\rightsquigarrow$) to the language with the following semantics:

Definition 6 (Semantics for '$\rightsquigarrow$')

(6) $\langle s_1, s_2\rangle \uparrow (\phi \rightsquigarrow \psi) =$
$$\langle\{w \in s_1 : (\langle s_1, s_2\rangle \uparrow^* \phi) \uparrow^* \psi = \langle s_1, s_2, \rangle \uparrow^* \phi\}, s_2\rangle;$$

(vi) $\langle s_1, s_2\rangle \uparrow^* (\phi \rightsquigarrow \psi) = (\langle s_1, s_2\rangle \uparrow (\phi \rightsquigarrow \psi))^*$.

Quite like the original Ramsey conditional, $\ulcorner\phi \rightsquigarrow \psi\urcorner$ invites the agent to run a test on his information state: If learning that ϕ and then learning that ψ would add no more information than learning that ϕ, then the information state passes the test. Otherwise, it fails. Equivalently: If learning that ϕ commits the agent to ψ, then update with $\ulcorner\phi \rightsquigarrow \psi\urcorner$ returns the original information state; otherwise, the result is the absurd state. The crucial difference between $\ulcorner\phi \Rightarrow \psi\urcorner$ and $\ulcorner\phi \rightsquigarrow \psi\urcorner$ then comes down to this: In order to entertain the antecedent of the former, the agent only needs to assume that ϕ is true; in order to entertain the antecedent of the latter, the agent needs to adopt the information state he would be in, were he to genuinely accept that ϕ is the case, i.e. he needs to adopt a potential belief state.

The Deduction Theorem is valid for '$\rightsquigarrow$', i.e. it is guaranteed that for all $\langle s_1, s_2\rangle$ and any sentences ϕ, ψ in $\mathcal{L}^{\mathsf{B}}_{\Rightarrow}$ extended by the binary connective '$\rightsquigarrow$', $\langle s_1, s_2\rangle \uparrow^* \phi \vDash \psi$ iff $\langle s_1, s_2\rangle \vDash \phi \rightsquigarrow \psi$. Accordingly, if we introduce the following rule of inference:

(MP*) $\ulcorner\phi \rightsquigarrow \psi\urcorner \vdash \ulcorner\phi \supset \psi\urcorner$

then the current framework predicts that rational commitment is closed under (MP*), i.e. that a rational agent who accepts $\ulcorner\phi \rightsquigarrow \psi\urcorner$ is committed to $\ulcorner\phi \supset \psi\urcorner$. And indeed, it is impossible to construct Thomason Conditionals for the connective '$\rightsquigarrow$'. Thus the following sound horrible if taken literally:

If I learn that Sally is deceiving me, I do not believe it (because she is so clever).

If I learn that Reagan works for the KGB, I'll never believe it.

If I learn that Reagan is bald, no one outside his immediate family knows it.

The counterexample to (MP) presented in §3 is is based on the possibility of constructing Thomason Conditionals for '$\Rightarrow$'. Since it is impossible to construct Thomason Conditionals for '$\rightsquigarrow$', there is, as far as I can see, no corresponding counterexample to (MP*). Accordingly, it seems safe to say that (MP*) is a valid rule of inference and that rational commitment is closed under (MP*). The fact that the current framework can offer something in the close neighbourhood of the intuitions articulated by Gillies suggests very strongly that it is on the right track and that the distinction between supposing and accepting is of outstanding importance for a successful theory of conditionals.

5.2 Comparison with *Must* and *Might*

It has long been observed that epistemic modals are similar to attitude verbs in their tendency to give rise to Moorean paradoxical phenomena. That is, the following construction sounds as bad as the original Moorean construction:

(a) # It is raining, and it might not be raining.

On the other hand, Yalcin (2007) observes that epistemic modals and attitude verbs behave differently once embedded in conditional constructions. The original Moorean construction can serve as the antecedent of a conditional, while the construction in (a) cannot:

(b) If it is raining and I do not believe that it is raining, then there is a truth I do not believe.
(c) # If it is raining and it might not be raining, then there is a truth I do not believe.

We can explain this contrast as follows. What needs to be assumed is that while belief self-ascriptions test on the second position of an information state, epistemic modals run tests on the first position:

Definition 7 (Semantics for 'Must' and 'Might')
(7*) $\langle s_1, s_2\rangle \uparrow \Box\phi = \langle\{w \in s_1 : \langle s_1, s_2\rangle \uparrow \phi = \langle s_1, s_2\rangle\}, s_2\rangle$;
(vii) $\langle s_1, s_2\rangle \uparrow^* \Box\phi = (\langle s_1, s_2\rangle \uparrow \Box\phi)^*$;
(8*) $\Diamond\phi =_{def} \neg\Box\neg\phi$.

Consider an acceptance base $\langle s_1, s_2\rangle$: The assumption that ϕ guarantees that the information carried by ϕ is now among the information encoded in the first position. Since $\ulcorner\Diamond\neg\phi\urcorner$ is a test on the first position, the conjunction $\ulcorner\phi \wedge \Diamond\neg\phi\urcorner$ leads to the absurd state even if it occurs as the antecedent of a conditional. In contrast, $\ulcorner\phi \wedge \textsc{Bel}(\neg\phi)\urcorner$ is consistent as long as it occurs as the antecedent of the conditional. In other words, $\ulcorner\phi \wedge \textsc{Bel}(\neg\phi)\urcorner$ can be coherently assumed, while $\ulcorner\phi \wedge \Diamond\neg\phi\urcorner$ cannot. This explains why (b) sounds all right, while (a) as well as (c) are unacceptable.

Besides accounting for the intuitions worked out above, this semantics for epistemic modals has the attractive feature that it captures the central logical aspects of epistemic modals. For example, consider the following valid line of reasoning:

(d) Jones is in Paris or Jones is in Rome. So, if he is not in Paris, then he must be in Rome; and if he is not in Rome, he must be in Paris.

The suggested semantics for epistemic modals predicts that this line of reasoning is valid since it yields the following:

(e) $p \vee q \models \neg p \Rightarrow \Box q$
(f) $p \vee q \models \neg q \Rightarrow \Box p$

Updating with $p \vee q$ yields an information state in which every world in s_1 and s_2 is a p-world or a q-world. If I now assume that p is false, then the only worlds remaining in s_1 are q-worlds, and so I have to hypothetically accept that q must be true. On the other hand, eliminating all q-worlds in s_1 by assuming that q is false requires me to hypothetically accept that p must be true.[15]

I conclude that the current framework captures the main logical features of the epistemic modals expressed by 'must' and 'might'. In particular, it allows to reconstruct which logical features these epistemic modals share with attitude ascriptions and where they disagree.

5.3 Revision

Assume an agent with regular information state who believes that ϕ. Such an agent has reason to accept the conditional $\ulcorner \text{Bel}\ (\phi) \Rightarrow \phi \urcorner$. But conditionals of this form must not turn out to be supported by every regular information state: Assuming that S does not believe in the truth of ϕ, S has no reason to assume the truth of ϕ under the condition that S believes that ϕ – unless, of course, S takes himself to be infallible. Strengthened by some reasonable assumptions, this is the prediction made by the framework developed here. Let me explain.

Indicative conditionals presuppose that adding the antecedent to the information state does not lead to the absurd state. If this presupposition is not fulfilled, accommodation is in order: If the agent is cooperative, he will temporally revise his information so that the presupposition is fulfilled. In many cases, this demands that information is temporarily retracted from the information state. Retraction, however, is not always the most efficient method of processing a non-open conditional. Sometimes it is sufficient to add information to the information state in order to avoid the absurd state by updating with the antecedent. Take a conditional the antecedent of which is itself a conditional, i.e. a conditional of the form $\ulcorner(\phi \Rightarrow \psi) \Rightarrow \chi\urcorner$. Now suppose that the information state of the evaluating agent is such that $\langle s_1, s_2\rangle \uparrow (\phi \Rightarrow \psi) = \langle \varnothing, s_2\rangle$. In that case, the presupposition of the indicative conditional that assuming the antecedent does not lead to the absurd state is violated. What should the agent do if he is cooperative? A natural reaction of the agent is that he modifies his information state by assuming the appropriate conjunction of the form $\ulcorner\neg(\phi \wedge \neg\psi)\urcorner$. As long as this modification itself does not lead to the absurd state, the agent

[15] In line with what was said before, the assumption that Jones is not in Paris (Rome) does not require me to assume that I believe that Jones is in Rome (Paris). Again, that is the right thing to say, since neither 'If Jones is not in Paris, I believe he is in Rome' nor 'If Jones is not in Rome, I believe he is in Paris' are constructions I should be committed to accept.

can now assume ⌜$\phi \Rightarrow \psi$⌝ without ending up in the absurd state, and check whether he hypothetically accepts χ on that supposition.

Assume now that both ϕ and ⌜$\neg\phi$⌝ are epistemic possibilities for an agent S, and that S considers a conditional the antecedent of which is of the form ⌜Bel (ϕ)⌝. As it stands, S's information state fails the test posed by ⌜Bel (ϕ)⌝, so revision is in order. The rational thing to do for S, I submit, is to modify the second position of the information state in such a way that the test is passed. This can be done without retraction: All S has to do is to eliminate ⌜$\neg\phi$⌝ as an epistemic possibility in s_2. We can describe this procedure in more general terms: If $\langle s_1, s_2\rangle$ is S's information state, then all S has to do is to find the maximal subset s'_2 of s_2 such that $\langle s_1, s'_2\rangle \uparrow \text{Bel}\ (\phi) = \langle s_1, s'_2\rangle$. This motivates the following:

Definition 8 (Tame Acceptance Bases) Consider arbitrary $\langle s_1, s_2\rangle \in I$ with non-empty first and second position, and assume further that $\langle s_1, s_2\rangle \uparrow \text{Bel}\ (\phi) = \langle \varnothing, s_2\rangle$. Then $\langle s_1, s_2\rangle$ is *tame* with respect to ⌜Bel (ϕ)⌝ iff, there is $s'_2 \subset s_2$ such that $\langle s_1, s'_2\rangle \uparrow \text{Bel}\ (\phi) = \langle s_1, s'_2\rangle$.

The following should be now an obvious modification of the semantics for conditionals:

Definition 9 (Updates on Tame Acceptance Bases) Assume that $\langle s_1, s_2\rangle$ is tame with respect to ⌜Bel (ϕ)⌝. Then we say that $\langle s_1, s_2\rangle \uparrow (\text{Bel}\ (\phi) \Rightarrow \psi)$ $= \langle\{w \in s_1 : (\langle s_1, s'_2\rangle \uparrow \text{Bel}\ (\phi)) \uparrow \psi = \langle s_1, s'_2\rangle \uparrow \text{Bel}\ (\phi)\}, s_2\rangle$, where s'_2 is the maximal subset of s_2 such that $\langle s_1, s'_2\rangle \uparrow \text{Bel}\ (\phi) = \langle s_1, s'_2\rangle$.

Updating with non-open conditionals the antecedent of which is a belief ascription thus involves the following two-stage process: First, the information state gets modified in such a way that it passes the test induced by the antecedent. Second, it is checked whether hypothetical update with the consequent adds any additional information. If not, the test posed by the conditional counts as passed; otherwise, the information state fails the test.[16]

Strengthened by the last two definitions, the double-indexed framework for conditionals predicts that there are sentences of the form ⌜Bel (ϕ) $\Rightarrow \phi$⌝ which are not supported by every information state: Acceptance bases which are tame with respect to ⌜Bel (ϕ)⌝ are not guaranteed to support such formulas, even if they are regular. The reason should be obvious: The modifications of the information state which allow it to pass the test encoded in the antecedent only affect the second position. The following update with ϕ affects the first position

[16]This suggestion can be extended to conditionals of the form ⌜$(\phi \Rightarrow \psi) \Rightarrow \chi$⌝ as follows: $\langle s_1, s_2\rangle$ is *tame* with respect to ⌜$\phi \Rightarrow \psi$⌝ iff, there is a $s'_1 \subset s_1$ such that $\langle s'_1, s_2\rangle \uparrow (\phi \Rightarrow \psi) = \langle s'_1, s_2\rangle$. If $\langle s_1, s_2\rangle$ is tame with respect to ⌜$\phi \Rightarrow \psi$⌝, then $\langle s_1, s_2\rangle \uparrow (\phi \Rightarrow \psi) \Rightarrow \chi =$ $\langle\{w \in s_1 : \langle s'_1, s_2\rangle \uparrow (\phi \Rightarrow \psi) \uparrow \chi = \langle s'_1, s_2\rangle \uparrow (\phi \Rightarrow \psi)\}, s_2\rangle$ where s'_1 is the maximal subset of s_1 such that $\langle s'_1, s_2\rangle \uparrow (\phi \Rightarrow \psi) = \langle s'_1, s_2\rangle$. This treatment is similar in spirit to Levi's (1988) treatment of iterated conditionals, but has the additional advantage of being motivated by considerations about presupposition and accommodation.

and might thus very well make a difference to the information state. This is the intuitive result: Hypothetically assuming that I believe that ϕ should not force me assume that ϕ is true. A double-indexed account for conditionals allows us to capture this intuition in a formal framework.

Ackknowledgements

I would like to thank audiences at the University of Texas at Austin, at the University of Southern California, and at the 36^{th} annual meeting of the Society for Exact Philosophy at the University of Wyoming for discussion and comments. Special thanks to an anonymous *Synthese* referee and to Josh Dever for their very useful comments. The usual disclaimers apply.

References

Arló Costa, H. and Levi, I. (1986). "Two Notions of Epistemic Validity". *Synthese*, 109 (2): 217–262.

Beaver, D. I. (2001). *Presupposition and Assertion in Dynamic Semantics*. CSLI Publications, Stanford, CA.

Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford.

Chalmers., D. J. and Hájek, A. (2007). "Ramsey + Moore = God". *Analysis*, 67 (2): 170–172.

van der Does, J., Groeneveld, W., and Veltman, F. (1997). "An Update on Might". *Journal of Logic, Language and Information*, 6 (4): 361–380.

Edgington, D. (1995). "On Conditionals". *Mind*, 104 (414): 235–329.

van Fraassen, B. (1980). "Review of Brian Ellis, *Rational Belief Systems*". *Canadian Journal of Philosophy*, 10 (3): 497–511.

Fuhrmann, A. (1989). "Reflective Modalities and Theory Change". *Synthese*, 81 (1): 115–134.

Gärdenfors, P. (1986). "Belief Revisions and The Ramsey Test for Conditionals". *Philosophical Review*, 95 (1): 81–93.

Gillies, A. (2001). "A New Solution to Moore's Paradox". *Philosophical Studies*, 105 (3): 237–250.

Gillies, A. (2004a). "Epistemic Conditionals and Conditional Epistemics". *Noûs*, 38 (4): 585–616.

Gillies, A. (2004b). "New Foundations for Epistemic Change". *Synthese*, 138 (1): 1–48.

Gillies, A. (2006). "What Might be the Case after a Change in View". *Journal of Philosophical Logic*, 35 (2): 117–145.

Groenendijk, J., Stokhof, M., and Veltman, F. (1996). "Coreference and Modality". In Lappin, S., editor, *Handbook of Contemporary Semantic Theory*, pages 179–214. Blackwell, Oxford.

Hansson, S. O. (1999). *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Press, Dordrecht.

Jackson, F. (1987). *Conditionals*. Basil Blackwell, Oxford.

Levi, I. (1977). 'Subjunctives, Dispositions and Chances". *Synthese*, 34 (4): 423–455.

Levi, I. (1980). *The Enterprise of Knowledge*. MIT Press, Cambridge, MA.

Levi, I. (1988). "iteration of conditionals and the ramsey test". *Synthese*, 76 (1): 49–81.

Levi, I. (1996). *For the Sake of the Argument*. Cambridge University Press, New York.

Lewis, D. K. (1986). "Postscripts to Probability of Conditionals and Conditional Probability". In Lewis, D., editor, *Philosophical Papers*, volume 2, pages 152–156. Oxford University Press, New York.

Lowe, E. (1987). "Not a Counterexample to Modus Ponens". *Analysis*, 47 (1): 44–47.

McGee, V. (1985). "A Counterexample to Modus Ponens". *The Journal of Philosophy*, 82 (9): 462–471.

Ramsey, F. P. (1990). "General Propositions and Causality". In Mellor, D. H., editor, *Philosophical Papers*, pages 145–163. Cambridge University Press, Cambridge.

Rott, H. (2001). *Change, Choice and Inference*. Oxford University Press, New York.

Shoemaker, S. (1995). "Moore's Paradox and Self-Knowledge". *Philosophical Studies*, 77 (2-3): 211–228.

Sinnott-Armstrong, W., Moor, J., and Fogelin, R. (1986). "A Defense of Modus Ponens". *The Journal of Philosophy*, 83 (5): 296–300.

Veltman, F. (1996). "Defaults in Update Semantics". *Journal of Philosophical Logic*, 25 (3): 221–261.

Woods, M. (1997). *Conditionals*. Clarendon Press, Oxford. Edited by David Wiggins, with a commentary by Dorothy Edgington.

Yalcin, S. (2007). "Epistemic Modals". *Mind*, 16 (464): 983–1026.