Belief revision conditionals: *basic* iterated systems

Horacio Arló-Costa [1]

Carnegie Mellon University, Philosophy,
Pittsburgh, PA 15213, USA
hcosta@andrew.cmu.edu

Abstract

It is now well known that, on pain of triviality, the probability of a conditional cannot be identified with the corresponding conditional probability [25]. This surprising impossibility result has a qualitative counterpart. In fact, Peter Gärdenfors showed in [13] that believing 'If A then B' cannot be equated with the act of believing B on the supposition that A – as long as supposing obeys minimal Bayesian constraints.

Recent work has shown that in spite of these negative results, the question 'how to accept a conditional?' has a clear answer. Even if conditionals are not truth-carriers, they do have precise acceptability conditions. Nevertheless most epistemic models of conditionals do not provide acceptance conditions for iterated conditionals. One of the main goals of this essay is to provide a comprehensive account of the notion of epistemic conditionality covering all forms of iteration.

First we propose an account of the *basic* idea of epistemic conditionality, by studying the conditionals validated by epistemic models where iteration is permitted but not constrained by special axioms. Our modelling does not presuppose that epistemic states should be represented by belief sets (we only assume that to each epistemic state corresponds an associated belief state). A full encoding of the basic epistemic conditionals (encompassing all forms of iteration) is presented and a representation result is proved.

In the second part of the essay we argue that the notion of change involved in the evaluation of conditionals is *suppositional*, and that such notion should be distinguished from the notion of *updating* (modelled by AGM and other methods). We conclude by considering how some of the recent modellings of iterated change fare as methods for iterated supposing.

[1] Special thanks to Rohit Parikh for coments on previous drafts of this paper.

1 Introduction

(RT) 'If A, then B' is accepted with respect to an epistemic state K iff B is accepted in the minimal revision of K needed to accept A for the sake of the argument.

RT is a neat and elegant idea. [2] How to make it precise has been, nevertheless, the object of some controversy. The terms that are left vague in RT are: 'epistemic state', 'hypothetical revision', and 'acceptance'. Different implementations of RT typically disambiguate these terms in different ways. I will start by offering a review of some of the basic alternatives.

First consider epistemic states. One common maneuver is to represent them by closed sets of sentences (*theories* in logical parlance). For example, the AGM trio adopted this position. [3] Sentences in those *belief sets* represent the sentences that a certain agent is *committed* to believe at a certain time. This move is now common in computer science circles, but it is far from being common in philosophy, where epistemic states are often encoded as probability functions.

Richer qualitative encodings of states have been proposed by many authors. Wolfgang Spohn, for example, has advocated in several papers the use of *well ordered partitions* (WOP). Epistemic states can also be represented by *Grove orderings* (which encode more structure than WOPs), or by *ordinal conditional functions* (which encode more structure than the two previous modellings). [4] .

No substantial assumptions about the structure of epistemic states will be made in this paper, aside from the fact that the belief states associated with them are belief sets (i.e. closed sets of sentences in a non-modal language). So, to each epistemic state E correspond a belief set ρ(E), representing what is believed by an agent in this epistemic state. [5]

---

[2] This idea is usually atributed to Frank Ramsey. Nevertheless Ramsey never endorsed (RT). For historical and conceptual details related to this issue see [2].

[3] See [13] for an introduction to AGM.

[4] This way of representing states has recently gained some momentum. Some representative examples of this trend are the elaboration of Sphon's ideas recently developed by Darwiche and Pearl in [10] as well as the proposal presented by Hans Rott in [29]

[5] Since no further restrictions are imposed until the last section of the paper, we can have different epistemic states sharing identical belief sets. Many philosophers require the complete determination of epistemic states by their doxastic components. This requirement will be discussed in section 5 of this paper, but the main representation result proved in Theorem 1 is completely general and therefore it

If E is the current epistemic state and ρ(E) is its associated belief set, one can use the notation E*A to indicate the minimal perturbation of E needed to *suppose* A. Of course with the new state E*A will be associated its coresponding belief set ρ(E*A).

What about acceptance? Acceptance of simple Boolean formulae in an epistemic state E can now be represented by set theoretical membership in ρ(E). Peter Gärdenfors proposed, in addition, to represent acceptance of conditionals with respect to an epistemic state E, by set theoretical membership in its associated belief set ρ(E) [13]. These proposals identify acceptance of modal and non-modal formulae with belief in the truth of the propositions that they express. This idea, in turn, entails that modal and conditional sentences do express propositions and are truth value bearers.

(GRT) $A > B \in \rho(E)$ iff $B \in \rho(E^*A)$ [6]

The view of conditionality induced by GRT is now well known. For example, we now know that GRT cannot be complemented, on pain of triviality, with a test for negated conditionals (see [14]):

(GNRT) $\neg\,(A > B) \in \rho(E)$ iff $B \notin \rho(E^*A)$

If the idea of GRT is to provide acceptance conditions for truth-value bearing conditionals, such a limitation is not surprising. In fact, GRT and GNRT imply that $A > B$ is rejected if and only if $A > B$ is not accepted. In other words, the net effect of complementing GRT with GNRT is to rule out the possibility of being in suspense about a conditional. But, of course, GRT seems to have been conceived in order to allow for such possibility.

Peter Gärdenfors showed also in [13] that there are no non-trivial conditionals generated by GRT when * is AGM. Recent refinements of this result seem to indicate that GRT is incompatible with principles of belief change that Ramsey explicitly required in his test (see [2] for details). In particular, GRT is incompatible with the following intuitive condition:

(*Consistent expandability*) If neither A nor $\neg$A is in ρ(E), then $\rho(E^*A) = Cn(\rho(E) \cup \{A\})$, where Cn is a classical notion of logical consequence.

How to interpret these results? Ramsey's remark in 'General propositions and

---

does not depend on such constraint.

[6] GRT stands for ' Gärdenfors' Ramsey test'.

causality,' can be of some help at this juncture:

> Many sentences express cognitive attitudes without being propositions; and the difference between saying yes or no to them is not the difference between saying yes or no to a proposition.[28]

Two implications of Ramsey's important remark are worth detailed consideration. First Ramsey suggests that one can develop acceptance conditions both for sentences that express propositions, and for sentences that, without being propositions, are, nevertheless cognitive carriers. This, I think, is one of the main insights of the program of *probabilistic semantics*, inititated by Ernest Adams [1]. Although Adams thinks that conditionals are not truth-carriers, he nevertheless thinks that it is possible to provide a precise and self-contained theory of acceptance for them. I do agree with this view, and my ambition in this and other writings is to help to develop a self-contained theory of interpretation for conditionals and other modals, whose ontological status as truth-carriers is at least open to question.[7] From now on I will use the term 'epistemic' to refer to these conditionals.[8]

Secondly, Ramsey suggests that the form of a theory of acceptance for propositions (in this case conditional propositions) does not need to coincide with the form of a theory of acceptance for epistemic conditionals. This idea has been less appreciated by philosophers, logicians and linguists. But I think that Ramsey's remark is quite right, and I hope that some of the following results will contribute to clarify it. In the following sections I will refer to Ramsey's proposal as *Ramsey's insight*.

We already considered a precise formulation of Ramsey's test for conditionals. GRT obviously provides an epistemic test for acceptance of conditional sentences. But it also seems obvious that this particular formulation of the test commits us to considering these conditionals as truth value bearing sentences, capable of being the object of belief, disbelief or epistemic suspense. So, GRT seems to be best seen as delivering acceptance conditions for conditional propositions.

Conditionals of that sort have been thoroughly (and independently) investi-

---

[7] I would like to show also that such a theory can be developed without appealing to any probabilistic paraphernalia. Moreover, I would like to suggest that a purely qualitative account can improve the performance of the standard probabilistic theories. In fact, probabilistic theories are typically unable to deliver a satisfactory account of iteration.

[8] Epistemic conditionals can be seen as truth-valued propositions, and can therefore be identified with set of possible worlds, as long as the corresponding propositions are undestood contextually. Such a view has been defended by Lindstrom and Rabinowicz in [20].

gated with the tools of possible worlds semantics. In [2] it is shown that some of the standard systems of possible worlds conditionals can be reconstructed (without appealing to the notion of truth) by supplementing GRT with the axioms of a notion of change recently axiomatized by the computer scientists Alberto Mendelzon and Hirofumi Katsuno [17]. It is not transparent which is the intended interpretation for such a notion (usually called *update*), but some of the possible interpretations seem to have an ontological flavor (for example, update fails to satisfy Consistent Expandability). [9]

Therefore, we can conclude that if there are such things as truth-value bearing conditionals, [10] GRT delivers acceptance conditions for them. Both negative and positive reasons support this view. On the positive side one can count the fact that GRT can be used to provide acceptance conditions for the possible worlds semantics conditionals. On the negative side, one should count the impossibility results mentioned above. These results show the impossibility of reconciling GRT with any theory of change motivated on epistemic grounds.

But the ambition of the probabilistic program in semantics, and its more qualitative sequels, has been and still is to develop a self-contained pragmatics, completely neutral about the ontological status of conditionals. This is probably the minimal goal of this program. More ambitiously, the idea can be also expressed as follows: 'Even if conditionals (or some class of conditionals) lack truth values at all, they have precise acceptance conditions. The basic idea of an epistemic reconstruction of conditionality is to make these conditions explicit.' [11]

The basic building blocks of an epistemic semantics for conditionals (of the kind that we want to consider here) have been developed by J. L. Mackie, Isaac Levi [12] and Peter Gärdenfors [13]. I contributed to develop this tradition in a series of recent papers. [13] A book-length defense of this view has been recently

[9] Meek and Glymour have suggested in [27] a distinction between *conditioning* and *intervening* that can be used to make sense of update and to understand how the operation relates to AGM. Most of the examples used to motivate update in the database literature can be modeled in terms of what Meek and Glymour call an intervention: i.e. an AGM change not with a proposition E, but with a more complex proposition I which is an action to bring about E.

[10] Many philosophers and logicians have expressed doubts about such a possibility. Among others, we should count E. Adams, M. Dummett, D. Edgington, A. Gibbard, I. Levi, J.L. Mackie, etc.

[11] One should stress that Ramsey's idea in his seminal piece is certainly related to this second project, rather than to the project of providing acceptance conditions for truth value carriers.

[12] See [21] and [22].

[13] See [3]. Part of the theory has been developed in collaboration with Isaac Levi [2].

offered by Levi [24].

Some further considerations about GRT can help us to understand how the new proposal works. One of the main ideas implemented in GRT is that conditionals and Boolean sentences should be treated symmetrically. Since the conditionals studied by the theory are presystematically understood as propositions, they are potential objects of belief, and therefore the conditionals generated by the test are stored in the *belief set* ρ(E), together with the Boolean formulae that the agent in question believes at the time. But, once one abandons the idea that the test treats conditional propositions, and one focuses on studying the conditionals that merely express the *commitments for change* already encoded in E, such a move is no longer methodologically sound. One can say instead that each epistemic state has an associated belief set, and that it *supports* the conditionals generated via Ramsey bridges. But in saying so we are tacitly introducing subtle changes in Gärdenfors' formulation of Ramsey's test, in such a way that the phrase 'the conditionals generated via Ramsey bridges', has now a new meaning. If s(E) denotes the *support set* associated with E, the formulation of the test is now as follows (the tests only apply to consistent s(E) and *unnested conditionals* A > B, where both A and B are purely Boolean):

(SRT) $\mathrm{A} > \mathrm{B} \in \mathrm{s(E)}$ iff $\mathrm{B} \in \rho(\mathrm{E}^{*}\mathrm{A})$.

We should supplement the test with the following three stipulations, explaining the relations linking s(E) and ρ(E).

(1) If A is a purely Boolean formula, then if $\mathrm{A} \in \mathrm{s(E)}$, then $\mathrm{A} \in \rho(\mathrm{E})$

(2) s(E) is closed under logical consequence

(3) $\rho(\mathrm{E}) \subseteq \mathrm{s(E)}$

The new *stratified* formulation of the test is now compatible with

(SNRT) $\neg(\mathrm{A} > \mathrm{B}) \in \mathrm{s(E)}$ iff $\mathrm{B} \notin \rho(\mathrm{E}^{*}\mathrm{A})$.

In addition, the notion * that appears in the test is now consistently expandable, and the impossibility results are circumvented. Both tests have been used in order to study the logic of non-nested conditionals and non-monotonic notions of consequence [3]. But they share a limitation with their probabilistic relatives: they do not offer a theory of acceptance for nested conditionals. Some researchers have invoked this asymmetry to suggest that there is something wrong with the epistemic account. The argument goes as follows: 'we are obviously capable of handling iterated conditionals (of reasonable complexity).

Therefore any adequate theory of acceptance of conditionals has to be able to provide an idealized account of this ability. Since the epistemic approach seems unable to do so, we should have suspicions its general adequacy.'

Responses to this argument have traditionally relied on the denial of its main premise. Michael Dummett, for example, in [11] says that: 'we have hardly any use, in natural language, for conditional sentences... in which the antecedent is itself a conditional, and hence we cannot grasp (their) content ...'. Alan Gibbard argued similarly in [15]: on the one hand we have certain forms of iteration that we cannot decipher (like iteration to the left). On the other hand, we have other forms of iteration that we can understand, e.g. conditionals in consequents. Nevertheless, iterated conditionals of this kind are always equivalent to sentences without embedded conditionals. In fact, sentences of the form A > (B > C), should be understood, according to this view, as *abbreviations* of sentences of the shape: (A ∧ B) > C. This position has been somewhat prevalent among researchers working with probabilistic models of acceptance of conditionals (like [1]).

I do not think that this response to the realist challenge is fully satisfactory. It seems that by denying the existence of embedded conditionals we are seriously crippling a general theory of acceptance of epistemic conditionals. Any theory of conditionals has to be able to make sense of our capacity to evaluate a sentence like:

'If there is no oxygen in the room, then it is not true that if I scratch the match it will light', or

'If the light goes on if you press the switch, then the electrician has finished his job.'

1.1 Iterated extensions of the Ramsey test

The stratified versions of the Ramsey test avoid paradox by breaking a symmetry. While GRT treats both modal and non-modal formulae symmetrically, SRT fractures this symmetry by producing a more sophisticated picture of acceptance of conditionals.

One can see this fracture as a cautious treatment of a serious problem. Even if conditionals are not truth carriers, there is a precise manner in which conditional sentences are supported by epistemic states. If one wants to preserve caution, support should not be conflated with acceptance understood as belief in the truth of a conditional proposition. So, the epistemic states used in SRT are more complex than the ones used in GRT. While the ones used in GRT have only a doxastic component, the ones needed to formulate SRT have both a doxastic component (the ρ-part) and a support component (the s-part).

The apparent strength of GRT resides in its capacity to deal with iteration, which is formally reflected by the fact that the test can use ρ-functions in both sides of the test. As we argued above, one must pay a high price for this symmetry: either paradox or a change of theme.

Can one improve the stratified tests in order to allow for iteration? In this essay I would like to consider what I believe is an evident strategy: restoring symmetry at the level of the support functions. One can appeal to the following iterated modifications of the stratified tests.

(ISRT) A $>$ B $\in$ s(E) iff B $\in$ s(E*A).

(ISNRT) $\neg$(A $>$ B) $\in$ s(E) iff B $\notin$ s(E*A).

The tests only apply in the case in which s(E) is consistent and B is in $L_{>}$. The logic of iterated epistemic conditionals can then be studied as usual by appealing to *epistemic models* (EM). An epistemic model is a quadruple $\langle \mathbf{E}, \rho, s, * \rangle$ $= \mathcal{M}$ where * is a belief revision function obeying (ISRT) and (ISNRT), **E** is a set of epistemic states closed under revision, ρ is a belief function, mapping each member E of **E** to a belief set ρ(K), and s a support function, mapping *belief states* to the sets of all formulae that these states support. A conditional sentence C is valid in M if and only if C $\in$ s(K), for every K in E.

The main motivation of this article is to use the EMs in order to provide firm foundations for the study of iterated epistemic conditionals.

1.2 Basic *epistemic conditionals*

Models of iterated belief revision have recently proliferated. Typically the AGM framework is either extended or revised in order to accommodate special axioms that legislate on iterated revisions. Then, the Ramsey schema is used to study the conditional reflection of the special constraints on iteration.

Perhaps the most controversial models are the ones that allow for conditional updating. What should count as the minimal perturbation of the current epistemic state to 'suppose' A $>$ B? Here 'supposing' cannot be 'supposing true', because this will commit us to the idea that conditionals are, after all, truth-bearers. Perhaps 'supposing' should be equated in this case with certain kind of fantasizing according to which the current epistemic state is transformed into another state supporting the 'supposed' conditional. [14] It is less clear in what sense this transformation can be constructed as a *minimal* perturbation of the current epistemic state compatible with that shift. [15]

Even if one focuses only on conditionals iterated to the right, there is little agreement about the correct principles of iterated change that go with the Ramsey test. Most of the theories of iteration in the literature seem too domain-dependent to be considered as universal constraints on iterated supposing. Actually the very distinction between belief change and change for the sake of the argument is often overlooked, which makes the matter even more obscure.

So, our strategy in this paper will follow the cautious line proposed by Brian Chellas in his paper 'Basic conditional logic'. [16] In this paper Chellas did not try to provide a substantive theory about some particular kind of conditionality – counterfactual or deontic, for example – but to study a *basic* notion of conditionality common to all the substantive theories of truth-valued conditionals. Our aim in the following sections will be to study a *basic* notion of *epistemic* conditionality.

2 Basic epistemic conditionals

2.1 Logical preliminaries

Let L_0 be a language containing a complete set of Boolean connectives, including the *falsum* and *verum* constants $\perp$ and $\top$. The set of wffs of L_0 is defined in the usual manner. In addition to L_0 we consider an operation Cn that takes sets of sentences in L_0 to sets of sentences in L_0. Cn is assumed to be a *consequence operation*, i.e. satisfying the classical Tarskian properties, [17] as well as finiteness, [18] consistency, [19] and superclassicality. [20] In addition Cn is such

[14] We follow here Levi's idea in [24], page 71.
[15] Two (conflicting) models are provided by Hansson in [16] and Boutilier in [6]. Boutilier represents states as ordinal functions, and proposes to construct a conditional shift, as the minimal perturbation of the ordinal function needed to make the conditional acceptable.
[16] See [7].
[17] For any sets K and H, such that K $\subseteq L_0$ and H $\subseteq L_0$, K $\subseteq$ Cn(K), Cn(K) = Cn(Cn(K)), and if K $\subseteq$ H, then Cn(K) $\subseteq$ Cn(H).
[18] For all X $\subseteq L_0$, Cn(X) = $\cup${C(X'): X' is a finite subset of X}.
[19] $\perp \notin$ Cn($\emptyset$).
[20] Cn($\emptyset$) contains all substitution instances of classical tautologies.

that for all A, B $\in L_0$ and K $\subseteq L_0$; A $\rightarrow$ B belongs to Cn(K) iff B ϵ Cn(K $\cup$ {A}).

Belief states of rational agents will be represented by belief sets satisfying the following rational constraints:

A *belief set* is a subset K of L_0 which satisfies the following conditions: (1) K is non-empty, (2) if A $\in$ K and B $\in$ K, then (A $\wedge$ B) $\in$ K, and (3) if A $\in$ K, and A $\rightarrow$ B $\in$ Cn($\emptyset$), then B $\in$ K.

Let T_{L_0} be the set of all belief sets that can be expressed in L_0. An *expansion* of a belief set K with a sentence A $\in L_0$ is the set Cn(K $\cup$ {A}). K $\in T_{L_0}$ is *consistent* if for some A $\in L_0$, A $\notin$ K.

Nothing substantial will be assumed about *epistemic states* E aside from the fact that their associated belief states ρ(E) are belief sets. An epistemic state E will be called *consistent* iff its associated belief set ρ(E) is consistent.

Let $L_>$ be the smallest language such that: (1) $L_0 \subseteq L_>$, (2) if A, B $\in L_>$, then A $>$ B $\in L_>$ and (3) $L_>$ is closed under the Boolean connectives. Let C be an operation that takes sets of sentences in $L_>$ to sets of sentences in $L_>$. C is assumed to be a *consequence operation*, i.e. satisfying the classical Tarskian properties, as well as finiteness, consistency and superclassicality. Of course, Cn $\subseteq$ C. In addition C is such that for all A, B $\in L_>$ and K $\subseteq L_>$; A $\rightarrow$ B belongs to C(K) iff B $\in$ C(K $\cup$ {A}).

A *conditional support set* is a subset K of $L_>$ which satisfies the following conditions: (1) K is non-empty, (2) if A $\in$ K and B $\in$ K, then (A $\wedge$ B) $\in$ K, and (3) if A $\in$ K, and A $\rightarrow$ B $\in$ C($\emptyset$), then B $\in$ K.

$T_{L_>}$ is the set of all conditional support sets. S $\in T_{L_>}$ is *consistent* if for some A $\in L_>$, A $\notin$ S.

2.2 Basic Epistemic Models

An Epistemic Model (EM) is quadruple $\langle \mathbf{E}, \rho, s, * \rangle$, where **E** is a set (heuristically: a ser of epistemic states); ρ is a function ρ: $\mathbf{E} \rightarrow T_{L_0}$; s is a function s: $\mathbf{E} \rightarrow T_{L_>}$; and * is a function *: $\mathbf{E} \times L_> \rightarrow \mathbf{E}$. **E** is closed under revisions and $\mathbf{B} = \mathrm{Rng}(\rho)$ is closed under expansions. The functions ρ, s, and * obey the following two constraints as well as IRT and INRT:

(c1) If A $\in L_0$ and A $\in$ s(E), then A $\in \rho$(E).
(c2) ρ(E) $\subseteq$ s(E).
(IRT) (A $>$ B) $\in$ s(E) iff B $\in$ s(E*A), where E is consistent.

(INRT) $\neg$ (A > B) $\in$ s(E) iff B $\notin$ s(E*A), where where E is consistent.

For every A $\in L_{>}$ and every $\mathcal{M} = \langle \mathbf{E}, \rho, s, * \rangle$, A is *satisfiable* in $\mathcal{M}$ if there is a consistent E $\in$ **E**, such that A $\in$ s(E). A is valid in $\mathcal{M}$ if A $\in$ s(E) for every consistent E $\in$ **E**. A is valid in a set of models S iff for every model $\mathcal{M}$ in S, A is valid in $\mathcal{M}$. A is *valid* if it is valid in all models. Finally A is *epistemically entailed* by B in $\mathcal{M} = \langle \mathbf{E}, \rho, s, * \rangle$, iff for every E in **E**, such that A $\in$ s(E), B $\in$ s(E).

Notice that the following constraint is derivable from the previous definition of *conditional support set*, given that C is finite (e.g. compact).

(c3) s(E) = C(s(E))

2.3 Syntax

First we need to define a conditional language smaller than $L_{>}$. Let $\mathcal{BC}$ be the smallest language such that if A, B $\in L_0$, C, D $\in BC$, then A > B, C > D, $\neg$ C, C $\wedge$ D $\in \mathcal{BC}$.

Consider now the conditional system $\mathcal{ECM}$. $\mathcal{ECM}$ is the smallest set of formulae in the language $L_{>}$ which is closed under (RCM) and (M), and which contains all instances of the axioms (I), (CC), (F) and all classical tautologies and their substitution instances in the language $L_{>}$.

I $A > \top$
CC $((A > B) \wedge (A > C)) \rightarrow (A > (B \wedge C))$
F $\neg$ (A > C) $\leftrightarrow$ (A > $\neg$ C), where C $\in \mathcal{BC}$.
M Modus ponens.
RCM If $\vdash B \rightarrow C$ then $\vdash (A > B) \rightarrow (A > C)$

The following completeness result shows the coincidence of the theorems of the system $\mathcal{ECM}$ and the conditionals validated by the EMs.

Theorem 1 *A $L_{>}$ formula A is valid iff A is a theorem in $\mathcal{ECM}$*

See appendix A for the proof.

3 Belief Revision Models and Elemental Epistemic Models

Axiom F encodes a property of basic epistemic models that we can call *fullness*. Fullness establishes that rational epistemic agents cannot be in suspense about conditionals – they cannot reject both a conditional and its negation. So, fullness should be rejected in any modelling of of truth value bearing conditionals. In fact, such models treat conditionals as any other truth value carrier, for which suspense should be allowed. [21] But, fullness is a typical property of epistemic conditionals. [22] In fact, it is a mandatory property of epistemic conditionals like the ones studied in the previous section, whose negations are understood in terms of lack of acceptance.

Fullness shoud not be confused with the so-called principle of *conditional excluded middle*:

(CEM) $(A > B) \vee (A > \neg B)$

This axiom fails to be validated in our system. Although the agents modelled by our theory are logically omniscient and fully introspective, the theory does not require the *complete* determination of the hypothetical state $\rho(E^*A)$ prompted by supposing A. Such a state can perfectly well be a (hypothetical) state of partial information. Fullness only reflects at the conditional level the fact that * is taken to be a universal *response sheme* to all possible suppositional items.

Axiom F mirrors syntactically the fact that the EMs implement both a test for acceptance (IRT) and a test for *rejection* of conditionals (INRT). As we said in passing in previous sections, one of the advantages of the EMs is that they allow for the implementation of both tests (in contrast, GRT and its negative counterpart are, on pain on triviality, inconsistent). Some philosophers have argued in favor of renouncing this theoretical advantage (see for example Collins' treatment of indicatives in [8]). Their argument can perhaps be reconstructed as follows: 'There are some conditionals whose status as truth carriers is, at least, open to doubt. We think that IRT is a good test of acceptance for these conditionals. But we also hope that a theory of truth can be given for them. In other words, we see a theory of acceptance for these conditionals as a prolegomenon to the construction of this truth-theory (perhaps one can use similar strategies than the ones that have been used to go from subjective to

[21] Suspense here refers to an epistemic scenario where neither a conditional not its negation are accepted.

[22] One might question fullness on the grounds that agents might fail to be aware of their commitments for change. Nevertheless the theory offered here does not intend to describe the actual *performance* of rational agents, but to propose normative constraints on the behavior of rational agents.

objective probabilities). Therefore we do not want to assume from the outset any additional constraint like INRT. In fact, assuming INRT (together with IRT) will be tantamount to assume that one cannot suspend judgement about conditionals. And we want to live this option open.'

Such theoretical ideas can be accommodated in our framework by adopting *elemental epistemic models*. They are just basic models where INRT is dropped.

Let $\mathcal{CM}$ be the smallest set of formulae in the language $L_{>}$ which is closed under (RCM) and (M), and which contains all instances of the axioms (I), (CC) in $L_{>}$, as well as all classical tautologies and their substitution instances in the language $L_{>}$. It is easy to see that the formulae validated by the elemental EMs can be axiomatized by $\mathcal{CM}$.

Moreover, when we complement the elemental models by:

(Reduction) For every epistemic model $\langle \mathbf{E}, \rho, s, * \rangle$, and every E in $\mathbf{E}$, s(E) $\subseteq$ ρ(E).

we get exactly Gärdenfors' Belief Revision Models. Gärdenfors' showed that the syntax validated by his models is (exactly) the one axiomatized by $\mathcal{CM}$ (see [13]).

3.1 Two types of conditionals?

Most contemporary theories of conditionals are bifurcated. Arguments for the existence of two types of conditionals can be presented concisely as follows: 'Acceptance conditions for one type of conditionals can be given in terms of IRT. Acceptance conditions for the other kind of conditionals can be given in terms of GRT, or what is equivalent, in terms of IRT plus Reduction. The latter type are clearly truth bearers, and their truth conditions can be independently delivered in terms of possible worlds semantics. It is not that clear that the conditionals in the former type are truth bearers, but we hope that IRT can be used as a heuristic tool to provide truth conditions that suit them.'

I will refrain from evaluating here the plausibility of this dualism. I would only like to remark that, even if one adopts a bifurcated theory, such a view seems to support what I called before *Ramsey's insight*. In fact, the previous results show that the form of a theory of acceptance for conditional propositions (either in the form of elemental models or in the form of elemental models obeying Reduction) differs from the form of a theory of acceptance for carriers of cognitive attitudes like the ones studied by the basic models.

The latent differences between the two theories are only reflected syntactically by languages that allow for iteration. In fact, notice that the axiom F is *essentially* iterated, in the sense that it has no non-nested instances. But one must keep in mind that the logics validated by the Belief Revision Models and the Basic Epistemic Models coincide for non-nested languages (see [3]).[23]

4 Extensions of the basic theory

In the following sections I will consider the pros and cons of extending the basic theory with the constraints of well known notions of belief change.

Due to its centrality in the theory of theory change, I will start with some important rationality postulates used by AGM.

(Inclusion) $\rho(\mathrm{E}^*\mathrm{A}) \subseteq \mathrm{Cn}(\rho(\mathrm{E}) \cup \{\mathrm{A}\})$

(Preservation) If $\neg$ A $\notin \rho(\mathrm{E})$, then $\mathrm{Cn}(\rho(\mathrm{E}) \cup \{\mathrm{A}\}) \subseteq \rho(\mathrm{E}^*\mathrm{A})$.

$\mathrm{Cn}(\rho(\mathrm{E}) \cup \{\mathrm{A}\})$ is usually abbreviated as $\rho(\mathrm{E}) + \mathrm{A}$, where $\rho(\mathrm{E}) + \mathrm{A}$ denotes the *expansion* of the belief set $\rho(\mathrm{E})$ with the sentence A. Inclusion and Preservation together entail that when a sentence A is consistent with the belief set $\rho(\mathrm{E})$, the belief set associated with $\mathrm{E}^*\mathrm{A}$ is just the expansion of $\rho(\mathrm{E})$ with A. A weaker condition has also been studied:

(Weak Preservation) If A $\in \rho(\mathrm{E})$, and $\rho(\mathrm{E})$ is consistent, then $\rho(\mathrm{E}) + \mathrm{A} \subseteq \rho(\mathrm{E}^*\mathrm{A})$.

There are two questions that one can ask regarding these postulates. One is purely formal: what are their conditional counterparts? The second is philosophically relevant: are the following postulates adequate to characterize the notion of iterated supposition involved in the evaluation of epistemic conditionals?

[23] One of the advantages of our account is that the differences are also made explicit syntactically for simpler languages, once the belief revision function used in the models is supplied with adequate constraints (see [3] and [2]).

4.1 Weak Preservation and Invariance

The rationality postulates presented in the previous section put some constraints on the dynamics of beliefs sets. If A is believed in ρ(E) and ρ(E) is consistent, the former postulates mandate ρ(E) = ρ(E*A), although E and E*A could be different. This is a trivial consequence of the fact that a belief set B can be associated with different epistemic states. Nevertheless some authors have suggested that such freedom should be curtailed. In fact, one can assume that there is one and only one epistemic state associated with a given belief set. Such a constraint can be introduced here via the following principle of *belief preeminence*: [24]

(BP) If ρ(E) = ρ(E'), then ρ(E*A) = ρ(E'*A).

BP has been assumed in some theories of suppositional or hypothetical change, like the one presented in [24]. [25] BP has also been assumed by other theories modeling 'real' change, but we will be mainly concerned here with the former application.

[24] The postulate has been used before under other names. For example, Friedman and Halpern call it *Propositionality* in [12].

[25] This statement needs some qualifications. The models of supposition studied in [24] are pairs $\langle B, M\rangle$, where B is a set of potential belief sets and M is a probability measure defined over all potential belief sets in B. Levi uses M to determine a measure of *informational value*. M is what Levi calls an *informational-value*-determining probability function. Independently of the details of Levi's construction what matters for us here is that for every model $\langle B, M\rangle$, expressions like (K*A)*B are always well defined in Levi's theory. In order to calculate (K*A)*B one starts with K. The fixity of the M function allows us to calculate K*A. Once one is in K*A, again by appealing to the fixed M-function, one can calculate the next revision. Revisions are not path-dependent. The result of revising a belief set K does not depend on its past history. It only depends on K and the M-function used in the model (and, of course, the particular way in which the M-function is used in order to construct the change operation *). The unconstrained version of the models considered in this paper behave in a very different manner. One can perfectly well have a pair of states E, E' and a sentence A, such that although ρ(E) = ρ(E'), ρ(E*A) is different from ρ(E'*A). For example, if an epistemic state E is modified by adding a sentence A compatible with ρ(E), Inclusion and Preservation will leave ρ(E) unchanged, but E might change to a different E'. But now, since * is a function that maps epistemic states and sentences to epistemic states, the outputs ρ(E*A) and ρ(E'*A) might differ – a concrete example will be considered later in this section. BP imposes a Markovian condition of path-independence that is a natural feature of the models used by Levi. It seems therefore that any reconstruction of Levi's models by appealing to a framework like the one that we are using in this paper, would require the imposition of BP.

Notice that Weak Preservation and Inclusion entail, in the presence of Belief Preeminence, the following constraint on iterated change

(Invariance) If A $\in \rho$(E), and ρ(E) is consistent, then ρ((E*A)*B) = ρ(E* B).

Our presentation of rationality constraints in this and the former section is not standard in the belief revision literature. AGM, for example, models belief change operations as mappings from belief sets and sentences to belief sets. The emphasis on this view is to provide one-step postulates which tell us what properties the next belief set ought to have, given the current belief set and the current evidence. If the interpretation is suppositional the idea is to provide postulates which tell us which properties a suppositional state ought to have given the current belief set and the current supposition. So, if we use B to denote belief sets, Inclusion and Preservation now look as follows:

(Inclusion) B $\subseteq$ Cn(B$\cup$ {A})

(Preservation) If $\neg$ A $\notin$ B, then Cn(B$\cup$ {A}) $\subseteq$ B*A.

When belief change postulates are presented in this form Invariance follows from them as a theorem:

(Invariance') If A $\in$ B, and B is consistent, then (B*A)*C = B*C.

Even when the idea of the AGM theory is to study one-step postulates, the theory does impose some constraints on iteration, like Invariance. The main goal of this section is to show that Invariance is not a good constraint on iterated supposing. The following example [26] will help us to make a case against Invariance:

Consider a spinner and a dial divided into three equal parts 1, 2, and 3. You (fully) believe that the spinner was started and landed in 1. Then you are asked to evaluate:

(I) If the spinner lands in an odd-numbered part, then if it is not 1, it is 3.

It seems pretty clear that you should say 'yes', i.e. that you should accept (I). But then it is also obvious that in this situation you should reject:

(II) If the spinner is not 1, then it is 3.

[26] The example is inspired by a similar one proposed in [24].

There are many examples of this kind. They were first brought to the attention of philosophers by Vann McGee in [26]. The examples are interesting because of their extreme generality. This is particularly valuable in an area where not even grammatical classification is uncontroversial. No one seems to agree about the correct classification of conditionals, and few phenomena seem to be valid across the different categories available in the literature. The McGee cases are one of these few phenomena. The impressive battery of examples exhibited in [26] is enough to convince any skeptic.

Different morals can be drawn from the examples. One particularly uncontroversial moral is to see the examples as showing that a rational agent can perfectly well believe a sentence A as true, and to accept (believe as true) a conditional of the form $A > (B > C)$, while he or she might find reasonable to reject the consequent $B > C$. This epistemic reading [27] of the McGee cases is all that is needed in order to mount a convincing case against the adequacy of Invariance, and therefore against Weak Preservation. [28] Notice that any invariant extended epistemic model obeys:

(EMP) $A, A > (B > C) \Rightarrow B > C$

Therefore no Invariant theory of iteration can make sense of the McGee cases. [29]

[27] 'Epistemic reading' has to be understood here in a very general manner, as applying both to our analysis and Gardenfors'.

[28] Van McGee seems to suggest in his essay that he wants to obtain a more ambitious conclusion from his examples. He indicates that his examples show that, if conditionals are truth value carriers, they do not obey in general the axiom $(A > B) \rightarrow (A \rightarrow B)$, where B is itself a conditional. Many objected against this way of interpreting the examples. Fortunately this strong diagnosis of the examples is not needed here. We only need the minimal epistemic base used by McGee in order to construct his more ambitious (and controversial) argument.

[29] A standard maneuver in order to dismiss McGee's examples is to say that – in these cases – the agent who accepts 'If A, then if B then C' does not really accept the iterated conditional $A > (B > C)$. What he does accept instead is a conditional with a conjunctive antecedent: $((A \wedge B) > C)$.

This observation can be interpreted in two possible manners. The first is usual in the literature on probabilistic semantics of conditionals: the conjunctive re-interpretation of 'If A, then if B then C' is just the symptom of a more generalized malaise: *there are no iterated conditionals at all.* This view is coherent, but it seems too draconian. It was briefly considered and dismissed in the introduction.

An anonymous referee reminded me of a second way of making sense of the rebuttal. There are, after all, iterated conditionals. $A > (B > C)$ is, under this point of view a legitimate candidate for acceptance. Then $A > (B > C)$ and its conjunctive re-interpretation are two syntactic objects used to express certain conditional beliefs. Now we are invited to consider that every time that an agent faces a McGee case he

What to give up? One option is to give up Belief Preeminence and to preserve Inclusion and Preservation. The following section will be devoted to analyze this option.

More radical conclusions can be obtained by introducing * as a mapping from belief sets and sentences to belief sets – as is customary in the AGM tradition. In this case Invariance is a direct consequence of Inclusion and Preservation. The immediate conclusion is that AGM (in its standard formulation) cannot be a model of supposition. One should adopt some weakening of the theory in order to avoid the unpalatable consequences induced by Invariance.

4.2 More about the adequacy of Weak Preservation: supposing vs. updating

Giving up Belief Preeminence is enough to block Invariance. Moreover, some of the theories that abandon BP, can handle McGee's cases rather well. An example is the theory recently proposed by Darwiche and Pearl in [10]. Darwiche and Pearl impose the following postulates on the dynamics of epistemic states in addition to Inclusion and Preservation:

Basic postulates

(Success) $A \in \rho(E^*A)$

(Consistency Preservation) $\rho(E^*A)$ is consistent if A is.

(Equivalence) If $E = E'$ and $A \leftrightarrow B$, then $\rho(E^*A) = \rho(E'^*A)$

chooses the sentence $A > (B > C)$ to express a conditional belief that as a matter of fact is not expressed by this sentence. What happens is that the agent entertains the conditional belief expressed by $((A \wedge B) > C)$. Moreover he really rejects the conditional belief expressed by the sentence used to reveal the conditional belief that he really has (i.e. he rejects the belief expressed by $A > (B > C)$, while he really accepts the conditional belief expressed by $((A \wedge B) > C)$)! This story is difficult to digest given the systematicity and generality of the McGee phenomenon. Why this systematic equivocation arises only in McGee cases? It seems much less problematic to say that the conditional belief expressed by the English conditional 'If A, then if B then C' is *always* the conditional belief expressed by the regimented sentence $((A \wedge B) > C)$. But we will see below that this *is* McGee's position. In fact, McGee thinks that we assert, believe or accept a conditional of the form $A > (B > C)$ whenever we are willing to assert, accept or believe the conditional $((A \wedge B) > C)$. Once one accepts the existence of iterated conditionals, the conjunctive move can only be part of an *explanation* of McGee cases, not an instrument used to dismiss them.

(Conjunctive revision) If ρ(E*A)+B is consistent, then ρ(E*A)+B = ρ(E*A $\wedge$ B)

In addition to these basic postulates Darwiche and Pearl impose the following set of special postulates for iterated revision:

C-postulates

(C1) If A entails B, then ρ((E*B)*A) = ρ(E*A)

(C2) If A entails $\neg$ B, then ρ((E*B)*A) = ρ(E*A) [30]

(C3) If B $\in$ ρ(E*A), then B $\in$ ρ((E*B)*A)

(C4) If $\neg$ B $\notin$ ρ(E*A), then $\neg$ B $\notin$ ρ((E*B)*A

Several useful examples are discussed in [10]. For example epistemic states can be encoded as *rankings* (or *ordinal conditional functions*). A ranking is a function κ from the set of all interpretations of the underlying language (worlds) into the class of ordinals. A ranking is extended to propositions by requiring that the rank of a proposition be the smallest rank assigned to a world that satisfies:

$\kappa(A) = min_{w \models A}\ \kappa(\mathrm{w})$.

The set of models corresponding to the *belief set* $\rho(\kappa)$ associated with a ranking κ is the set {w: k(w) = 0}. Darwiche and Pearl proved in [10] that the following method for updating rankings satisfy their postulates:

$$(\kappa \bullet A)(w) = \begin{cases} \kappa(w) - \kappa(A) & \text{if w} \models \text{A} \\ \kappa(w) + 1 & \text{otherwise} \end{cases}$$

It is clear that this method will handle the example of the spinner in a nice manner. The following picture shows how to encode the epistemic state of a person that just saw that the spinner landed in 1 – w1 is the 'world' representing the fact that the spinner landed in 1, and the same goes for 2 and 3.

[30] See [4] for comments on the general adequacy of this postulate.

κ	Possible worlds
0	w1
1	w2, w3
2	Rest of worlds

And the next two pictures show how the rankings evolve when updated by the proposition Odd = {w1, w3} first (κ_1), and then by $\neg$ 1 = { w2, w3} (κ_2).

κ_1	Possible worlds	κ_2	Possible worlds
0	w1	0	w3
1	w3	1	w2, w1
2	w2	2	Rest of worlds

While obviously, the result of directly updating κ with $\neg$ 1 yields:

κ_3	Possible worlds
0	w3, w2
1	w1
2	Rest of worlds

Supposing and updating are different kinds of mental acts. While the notion of supposing is clearly involved in the evaluation of conditionals, the notion of updating is involved in the actual process of changing view. The evaluation of a conditional like Odd $>$ ($\neg$ 1 $>$ 3) requires us to *suppose* Odd, not to actually change our minds in order to accommodate the information that the spinner landed in an odd number. The latter change might be modeled in many different ways. For example a body of recent literature on *non-prioritized* revision suggests that, depending on the structure of the initial state, not all revisions need to be successful.[31] But *supposing* Odd seems to require us to open a 'suppositional window' in which one focuses only on Odd states. This might be accomplished in many ways, none of which can be encoded via an

[31] Statistically inspired models like the method of *routine expansion* developed by Isaac Levi in [23] are even more liberal. Once an inconsistency is detected, one can question either the input, or the background or both.

invariant AGM method.[32]

Some notable exceptions notwithstanding the notions of supposing and updating have been usually conflated in the recent literature. This is surprising because the distinction between these two mental acts has considerable philosophical and formal pedigree. J.L. Mackie, for example, based his account of conditionals on the notion of supposing, which he carefully distinguished from the notion of learning. More recently supposing accounts of conditionals have been offered by Cross and Thomason, Collins, Levi, Skyrms and others.[33] Unfortunately there is no consensus about the properties that the notion of supposing has to have. There is more agreement when it comes to distinguish supposing from updating (both formally and conceptually).

We already saw that operations like Darwiche and Pearl's could be used to model some crucial aspects of supposing. Are they also good models of belief change? This is a controversial topic that we prefer not to touch here in detail. If one focus on Invariance, for example, the property is not universally imposed by all the recent theories of iterated change offered in the literature. For example, Lehmann's recent *revision* of belief revision [19] imposes Invariance axiomatically (even when the theory is formulated in a framework where the postulate could have been abandoned). On the other hand theories like Darwiche and Pearl's abandon BP and Invariance.

Consensus about the principles of belief change will be probably hard to obtain without anchoring the whole analysis on a broader perspective. For example, some authors have justified their views on the basis of an analysis of what are (or should be) the aims of inquiry (a Bayesian perspective has been defended in [23], and a reliability view has been offered in [18]). Alternatively one can rely on purely pragmatic considerations, rather than on what is in general required by rationality.

While one could argue about the adequacy of the principles of belief change, some of these doubts seem highly unmotivated when one focus on supposing. One might doubt whether update is successful or invariant. Such doubts seem bizarre in the case of supposing. Failure of Invariance and satisfaction of success seem *constitutive* of a reasonable notion of supposing.

We can recapitulate here our account of the role of Weak Preservation (WP). Let us concentrate first on a formulation of change rules as functions on epistemic states. In this case WP entails Invariance in the presence of Belief Preeminence and Inclusion. We just saw one way of blocking Invariance, via the

[32] If one models change by mapping belief sets and sentences into belief sets; one can directly conclude that no AGM method is capable of modeling supposing.

[33] The main references are [24], [9], [8] and [30].

rejection of Belief Preeminence.

This move is not available to researchers who impose Belief Preeminence, or that directly model change as functions on belief sets. In this case either WP or Inclusion should be given up. Some authors have argued in favor of abandoning WP, by pointing out that there are independent reasons for rejecting it. In fact, notice that the conditional reflection of Weak Preservation at the unnested level is:

(WP) $(A \wedge B) \rightarrow (A > B)$, for A, B $\in L_0$

Here is an example presented by Levi in [24] showing the inadequacy of (WP):

> **Example 2** (Levi): Suppose agent X is offered a gamble on the outcome of a toss of a fair coin where he wins \$ 1,000 if the coin lands heads and loses nothing if the coin lands tails. Let utility be linear in dollars. The expected value is \$ 500. X has to choose between this gamble and receiving \$ 700 for sure. X has foolishly (given his beliefs and values) accepted the gamble and won \$ 1,000. Y points out to him that his choice was foolish. X denies this. He says: 'If I had accepted the gamble, I would have won \$ 1,000.'

Against intuition, the acceptance of the last conditional is mandated by (WP). X believes that he accepted the gamble. He also believes that he won \$ 1,000. Therefore, according to (WP), X should accept the conditional in question. Levi has proposed in [24] a theory that rejects Weak Preservation. The theory accepts, nevertheless, Belief Preeminence.

There are further reasons not to abandon Inclusion, which we will consider in section 4.3. Roughly the idea is that one wants the conditional counterparts of the postulate.

4.2.1 Van McGee on supposing and the role of Consistency Preservation

In the previous sections we saw that the theory proposed by Darwiche and Pearl can be used to model certain forms of suppositional reasoning. We also considered a particular method of change satisfying Darwiche and Pearl's postulates and we saw how it can be used to deal with McGee cases.

McGee's own solution to his puzzling examples will be considered in this section. We will also show that the notion of supposition presupposed by McGee's theory cannot be modeled by any method obeying Darwiche and Pearl's axioms.

It is a familiar fact that the attitudes expressed by $A > (B > C)$ and its conjunctive re-interpretation ((A $\wedge$ B) > C) are sometimes difficult to disentangle. Researchers unconvinced by the theoretical plausibility of nested conditionals have traditionally appealed to this fact to deny that the acceptance of the English conditional 'If A, then if B then C' is ever related to the acceptance of any iterated conditional. ' If A, then if B then C' looks like an iterated conditional, but this is just an accident. The acceptance of ' If A, then if B then C' can always be explained in terms of the acceptance of ((A $\wedge$ B) > C).

McGee does accept the existence of nested conditionals. But he explains away the embedded conditional by postulating that A > (B > C) and its conjunctive reinterpretation really express the same attitude. In other words McGee postulates the validity of:

(Export-Import) $A > (B > C) \leftrightarrow ((A \wedge B) > C)$.

Then he proves an interesting result showing that accepting Export-Import has the consequence that either modus ponens fails for conditionals with conditional consequents *or* that the conditional ' $>'$ is the truth-functional conditional.

The epistemic counterpart of Export-Import is:

(Export*Import) $(E^*A)^*B = E^*(A \wedge B)$.[34]

Export-Import has fairly strong consequences. In the presence of Success the Export*Import postulate entails:

(Global Success) $A \in \rho((E^*A)^*B)$

In the presence of Success the Export*Import postulate also leads to violations of Consistency Preservation. In fact, consider the case when A is a propositional atom and B is $\neg$ A. Since Darwiche and Pearl's theory obeys Consistency Preservation, it is immediate that there is no Export*Import extension of their theory.

Preserving consistency is indeed one of the main goals of any theory modeling the process of updating epistemic states. But consistency preservation need not be an ideal governing all forms of suppositional reasoning. Adams, McGee and others tend to see supposition as a *cumulative* operation. According to this view (formally expressed by Global Success) the act of entertaining

[34] Or its more cautions counterpart: $\rho((E^*A)^*B) = \rho(E^*(A \wedge B))$.

(sequentially) two mutually contradictory suppositional inputs leads to a contradiction, not a case where the first supposition is retracted in virtue of the new one.[35] A formal model of a notion of supposition that strictly satisfies the Export*Import condition is offered in [5]. The gist of the proposal presented in [5] will be presented below after introducing an example. The example will motivate some of the issues discussed throughout the the section. It will also be used to illustrate the method used in [5].

Consider an spinner like the one used in the previous examples. Now the spinner has four numbers: 1, 2, 3 and 4. The initial epistemic state is:

κ	Possible worlds
0	w1, w4
1	w12, w3
2	Rest of worlds

Notice that the $\bullet$-rule sanctions the acceptability of the following conditional with respect to this epistemic state:

(1) If the winner lands in an Odd number, then if the number is different from 1, the number is either 3 or 4.

The rule also sanctions the acceptability of:

(2) If the winner lands in an Odd number different from 1, the number is 3.

Some philosophers and logicians[36] have found reasonable (both systematically and pre-systematically) to link rigidly the acceptance conditions of (2) and:

[35] This view has also been favored by researchers working in *inheritance networks*. See section 3 of [9].

[36] J. Collins, D. Edgington, F. Jackson, V. McGee and M. Woods should be counted among the theoreticians who have defended the Export*Import rule. Without plunging too deeply into Byzantine grammatical distinctions it should be said that only a subset of these authors think that the rule applies universally to all kind of conditionals. All of them think that it applies to some classes of conditionals. The scope of applicability varies according to the classification adopted. The two salient candidates are *indicatives* and *hypotheticals*. The conditionals used in our example fall under both categories.

(3) If the winner lands in an Odd number, then if the number is different from 1, the number is 3.

But the straight application of the •-rule recommends to accept (2), and (1) and to reject (3). The acceptance of (1) illustrates a failure of Global Success. The supposition Odd is not held fix throughout the evaluation of (1).

The basic idea behind Export*Import is that (2) and (3) express the same attitude. Notice also how different is (3) from:

(3') If I were to revise my epistemic state by Odd and then subsequently by $\neg$ 1, then I would believe that the winner is 3.

(3') should be rejected, but (3) should be accepted (and (1) rejected). The attitude of supposing that A is the case should not be conflated with the attitude of supposing that the current view has been updated after learning A. The second type of suppositional state is needed to evaluate (3'), while the first is needed to evaluate (3).[37]

Here is a different manner of making sense of the acceptability of (3). Say that the initial epistemic state is κ. Then supposing Odd can be modeled as the act of opening a 'suppositional window' that only includes Odd-options. This suppositional window should preserve as much as possible of the structure encoded in the initial ranking . For example it does not seem plausible to alter the ordering of worlds induced by κ. So, our first act of supposing will lead us to focus on:

κ'	Possible worlds
0	w1
1	w3
2	Rest of worlds

Now, when we suppose $\neg$ 1, with respect to κ', we get the desired result. This way of characterizing the notion of change involved in the evaluation of conditionals guarantees the property of Global Success suggested in the previous section. The items supposed at any point in a chain of suppositions are firmly maintained until the end of the process.

[37] Notice also that (3) and (3') seem to have different logical form. It is very difficult to turn (3') into a nested conditional. It seems more natural to see (3') as a simple (unnested) conditional with a temporal 'and then' connective in the antecedent.

The formal account presented in [5] does not appeal to rankings (we use well-ordered partitions or Grove systems). Nevertheless the operation can be defined in terms of rankings. We need a different characterization of rankings. In our previous characterization we defined rankings as mappings from the set of all interpretation of the underlying language into the class of ordinals. We will appeal now to a different characterization according to which rankings are understood as mappings from a set S of interpretations of the language to the class of ordinals. The ranking induced by the empty set of worlds is the inconsistent ranking κ_f.

If $S \cap A \neq \emptyset$ the rank assigned to A-worlds is determined in the usual way κ@A = κ(w) - κ(A). No $\neg$ A world is assigned a rank. If $S \cap A = \emptyset$, κ@A = κ_f.

Of course Consistency Preservation fails in this setting. Assume that S is nonempty and A is consistent. Say that our initial epistemic state is a ranking κ such that $\rho(\kappa)$ is non empty. When $S \cap A = \emptyset$, κ@A = κ_f, violating Consistency Preservation. Moreover, for every sentence B, (κ@A)@B = κ_f.

It is worth noticing that the example considered above illustrates a case where the suppositional inputs are mutually consistent. Informal analysis of conditionals defending the Export*Import rule tend to focus on cases of these kind. This might be justified by the empirical fact that nested English conditionals with contradictory antecedents are rare and hardly intelligible.

4.3 Inclusion

What about Inclusion? The natural unnested conditional counterpart of Inclusion is:

(CMP) $A > B \rightarrow (A \rightarrow B)$, with $A, B \in L_0$.

Does this run against McGee examples? The answer is no. The McGee cases show the inadequacy of any theory of supposition inducing the validity of a conditional axiom of the form CMP where B itself is a conditional. But CMP is a perfectly sound conditional axiom that one wants to preserve. The axioms that one does not want to have is:

(IMP) $(A > B) \rightarrow (A \rightarrow B)$, with $A \in L_0$, and $B \in L_{>} \backslash L_0$.

But it is easy to check that IMP is not validated by Inclusion in the context of our theory. Moreover in [3] is shown that there are no non-trivial epistemic models validating IMP. We will see in the following sections that the notion of acceptance characterized by the BRMs behaves in a very different manner.

4.4 Invariance and Inclusion in the context of the BRMs

The argument of this section can be summarized as follows: the BRMs deliver an inadequate theory of iteration, because every inclusive BRM also obeys EMP.

Notice that any notion of change compatible with (GRT) is monotonic:

(K*M) If $\rho(\mathrm{E}) \subseteq \rho(\mathrm{H})$, then $\rho(\mathrm{E}^*\mathrm{A}) \subseteq \rho(\mathrm{H}^*\mathrm{A})$.

Consider now the following weak version of Invariance:

$\rho((\mathrm{E}^*\mathrm{A})^*\mathrm{B}) \subseteq \rho(\mathrm{E}^*\mathrm{B})$, if $\mathrm{A} \in \rho(\mathrm{E})$, and $\rho(\mathrm{E})$ is consistent.

Remark 2 *Any monotonic and inclusive notion of change is weakly invariant.*

PROOF. Assume that $\mathrm{A} \in \rho(\mathrm{E})$, and $\rho(\mathrm{E})$ is consistent. By Inclusion: $\rho(\mathrm{E}^*\mathrm{A}) \subseteq \mathrm{Cn}(\rho(\mathrm{E}) \cup \{\mathrm{A}\}) = \rho(\mathrm{E})$. Therefore, by (K*M), $\rho((\mathrm{E}^*\mathrm{A})^*\mathrm{B}) \subseteq \rho(\mathrm{E}^*\mathrm{B})$ □

It is important to notice that all the theories of change that have been proposed in the literature in order to give epistemic models of the possible worlds conditionals obey (K*M). Good examples are [17] and [8]. These notions of change are qualitative versions of Lewis' *imaging*, a notion of change that imposes monotonicity at the qualitative level. So, the defenders of (GRT), and more generally the defenders of imaging, cannot give up (K*M). Then the only option is to give up Inclusion. But this is tantamount to giving up also:

(CMP) $\mathrm{A} > \mathrm{B} \rightarrow (\mathrm{A} \rightarrow \mathrm{B})$, with A, B$\in L_0$.

The problem is that there is no good reason to do so. Any defender of a notion of acceptance as 'believing true' should say that every rational agent who believes that both A and $\mathrm{A} > \mathrm{B}$ are true, should also believe that B is true. Therefore the account of acceptance of iterated conditionals using BRMs faces the following dilemma: either the models are constrained by Inclusion to have CMP, and then they are unable to deal with the McGee cases; or they are sensible to them, by relinquishing the unnested version of conditional modus ponens. This is a dilemma that all BRMs face, even the ones that are not weakly preservative – like the one developed in [8].

It is hard to see how to escape from such a dilemma. It seems that the BRMs face insurmountable problems to give an adequate analysis of iteration. According to many, the main defect of the epistemic account (using EMs) is its inability to give an adequate picture of iteration, while the BRMs are naturally suited to deal with iteration. Nevertheless, a more careful analysis of this idea shows that natural extensions of the stratified tests deliver a more adequate picture of iteration than the one given in terms of the BRMs.

5 Conclusion

The paper proposes a full treatment of iterated epistemic conditionals. A representation result is given for models that allow for iteration, but that do not prescribe especial constraints on the notion of supposition used in the models. We verified that even at this basic level, the theory departs from (deceptively) similar accounts, like the one done in terms of Belief Revision Models.

More substantial remarks about the notion of *supposition* involved in the evaluation of conditionals are offered in the last section. We concluded that a notion of iterated supposing (unlike the corresponding notion of iterated update) [38] cannot be Invariant, and we provided several arguments against Weak Preservation.This suggests that notions like AGM are not good candidates as an axiomatization of supposition. Other AGM axioms, like Inclusion, are adequate in the context of our theory. Surprisingly we verified, that every inclusive BRM is also invariant. This fact raises doubts about the adequacy of the BRMs for modelling iterated conditionals.

Acknowledgements

This work was stimulated considerably by many discussions with Isaac Levi on conditionals and the role of supposition. Many other discussions with Kevin Kelly have been also extremely helpful. In particular I am grateful for several formal suggestions that persuaded me about the possibility of proving a general completeness result covering *all* forms of iteration.

[38] Here 'update' refers to the actual process of changing view, not to the method of change proposed by Katsuno and Mendelzon.

I would like to thank also Matthias Hild, Joe Halpern, and Jesse Hughes for many helpful suggestions and comments on the material of this paper. Finally I would like to thank the valuable comments made by an anonymous referee.

A Proofs

Theorem 1: A $L_{>}$ formula A is valid iff A is a theorem in $\mathcal{ECM}$

PROOF.

Soundness: It is immediate to verify the validity of I, CC, and to check that the rules of inference preserve validity. The validity of F can be proved by induction on the complexity of formulae of $\mathcal{BC}$.

Completeness: To show the converse we will assume that a sentence α of $L_{>}$ is not a theorem in $\mathcal{ECM}$ and we will show that α is not supported by some epistemic state E of some epistemic model $\langle \mathbf{E}, \rho, s, * \rangle$. So, in the following we will construct explicitly an $\mathcal{ECM}$ model $\langle \mathbf{E}, \rho, s, * \rangle$ and we will exhibit an E $\in$ $\mathbf{E}$, such that α does not belong to s(E).

Some preliminary terminology: a sentence A is a *theorem* of a conditional logic L (usually written $\vdash_L$) A) just in case A $\in$ L. A is *derivable* in L from a set of sentences if and only if the set contains sentences $A_1, ..., A_n$ (n $\neq$ 0) such that and $\vdash_L$ (A_1 $\wedge$...$\wedge$ A_n) $\rightarrow$ A. A set of sentences in L is *consistent* just in case not every sentence is L-derivable from it.

Form now a list of all formulae of $L_{>}$ of the shape $\neg(A > B)$: g_1, g_2,..., g_n,... We suppose that each $L_{>}$-sentence of the shape $\neg$ (A > B) occurs at least once in this list. Now, with respect to this list, we construct an infinite sequence of sets:

I_0, I_1,..., I_n,...

as follows. As I_0 we take $\mathcal{ECM} \cup \{\neg\alpha\}$. Then, for each positive integer n we set:

$$I_{i+1} = \begin{cases} I_i, g_{i+1} & \text{if } I_i, g_{i+1} \text{ is consistent in } \mathcal{ECM} \\ I_i & \text{otherwise} \end{cases}$$

Set now: $\overline{Cn(\emptyset)}$ = C($\cup I_i$), where $\cup I_i$ denotes the union of all the infinitely many sets I_i.

The following series of sets of ordered pairs of sets can be constructed. We will adopt the following notation. Given an ordered pair K = $\langle B, C\rangle$, where B $\subseteq$ L_0 and C $\subseteq$ $L_>$: K_s = C, K_ρ = B.

$K_0 = \{\langle \overline{Cn(\emptyset)} \cap L_0, \overline{Cn(\emptyset)}\rangle\}$

$K_{n+1} = K_n \cup \{\langle\{C \in L_0 : A > C \in K_s\}, \{C \in L_> : A > C \in K_s\}\rangle$: K $\in K_n$, A $\in L_>\}$ $\cup$ $\{$ $\langle Cn(K_\rho \cup \{A\}), C(K_s \cup \{A\})\rangle$: K$\in K_n, A \in L_0\}$

We have now enough elements to define **E**, and the functions ρ, s and *.

(1) **E** $= \cup\, K_i$
(2) For all E $\in$ **E**, ρ(E) = E_ρ
(3) For all E $\in$ **E**, s(E) = E_s
(4) $\langle E_\rho, E_s\rangle$ * A = $\langle\{C \in L_0 : A > C \in E_s\}, \{C \in L_> : A > C \in E_s\}\rangle$

We should verify now that $\langle \mathbf{E}, \rho, s, *\rangle$ is an EM.

First we will check that for all E $\in$ **E**, s(E) is a consistent conditional support set containing all the theorems of $\mathcal{ECM}$, and that s(E) is complete with regard to formulae of the shape A > B $\in L_>$ (i.e. for all E $\in$ **E**, either A > B $\in$ s(E) or $\neg$(A > B) $\in$ s(E)). By construction the claim is true for K_0, which only contains the pair $\langle \overline{Cn(\emptyset)} \cap L_0, \overline{Cn(\emptyset)}\rangle$.

Now suppose that all pair of sets in K_n have these properties. Then we want to show that all pairs in K_{n+1} also have these properties.

We should check that for all E in K_{n+1}, s(E) = $\{\{$ $C \in L_> : A > C \in K_s$ $\}$: K $\in K_n$, A $\in L_>\}$ is a consistent conditional support set containing all theorems of $\mathcal{ECM}$, complete with respect of sentences of the shape A > B $\in L_>$.

First consistency. Assume by contradiction s(E) is inconsistent, for some E $\in$ K_{n+1}. Then, C, $\neg$ C $\in \mathcal{BC}$ belong to S(E). Therefore A > C, A > $\neg$ C are in K_s, for some K in K_n. Since by inductive hypothesis (IH from now on) K_s is a conditional support set containing all $\mathcal{ECM}$-theorems, $\neg$(A > C) $\in K_s$ (axiom F), against the assumed consistency of K_s.

Secondly we will show that s(E) is a conditional support set. It is easy to see that s(E) is non-empty (because, by (I), A > $\top$ $\in K_s$). By (CC) and IH, if A, B $\in L_>$ belong to s(E), A $\wedge$ B $\in$ s(E). Finally if C $\rightarrow$ B $\in$ C($\emptyset$) and C $\in$ s(E), since A > C $\in K_s$, by RCM (and IH) we have that A > B $\in K_s$, and therefore B $\in$ s(E).

Thirdly for completeness, assume that B $>$ C $\in L_>$, is such that B $>$ C $\notin$ s(E). Then A $>$ (B $>$ C) $\notin K_s$. Since K_s is complete, $\neg$ (A $>$ (B $>$ C)) $\in K_s$, and since all theorems of $\mathcal{ECM}$ are in K_s, A $> \neg$ (B $>$ C) $\in K_s$. Therefore $\neg$ (B $>$ C) $\in$ s(E).

Finally we should check that for all τ, such that $\tau \in \mathcal{ECM}$, $\tau \in$ s(E). Notice that $\top \to \tau \in$ C($\emptyset$) and that therefore (A $> \top) \to$ (A $> \tau$) is also an $\mathcal{ECM}$-thesis. By axiom (I), therefore, A $> \tau \in \mathcal{ECM}$. So A $> \tau \in K_s$, and $\tau \in$ s(E) as one wants.

It is easy to check that for all E in K_n, ρ(E) is a belief set (by using appropriate instances of axioms (I), (CC) and the rule of inference RCM).

It is also easy to check now that the Ramsey clauses are obeyed. If B $\in \rho$(E*A), with B $\in L_0$, E $\in$ **E**, then A $>$ B $\in$ s(E), and, since s(E) is consistent, $\neg$ (A $>$ B) $\notin$ s(E).

If A $>$ B $\in$ s(E), then B $\in \rho$(E*A), by definition. If $\neg$ (A $>$ B) $\in$ s(E), consistency requires A $>$ B $\notin$ s(E) and therefore B $\notin \rho$(E*A). Finally if B $\notin \rho$(E*A), A $>$ B $\notin$ s(E), and by completeness of s(E) with regard of formulae of this shape, $\neg$(A $>$ B) $\in$ s(E). An almost identical strategy is needed to check the case σ = s.

We need now to check that constraint (c1) is satisfied. This also will be done by induction. K_0 contains only the pair: $\langle\overline{Cn(\emptyset)} \cap L_0, \overline{Cn(\emptyset)}\rangle$, which satisfies (c1) by construction.

Assume now, by IH, that (c1) is satisfied by all K $\in K_n$. We will then prove that the constraint (c1) is satisfied by all the pairs E in K_{n+1}.

Our first target will be the expansions in EXP = { $\langle Cn(K_\rho \cup \{A\}), C(K_s \cup \{A\})\rangle$: K $\in K_n, A \in L_0$}. We need to check that for all B $\in L_0$, B is in Cn($K_\rho\cup${A}) whenever B is in C($K_s\cup${A}). So, assume that B is in C($K_s\cup${A}). Then (by assumed properties of C) A $\to$ B $\in$ C(K_s) = K_s. But then, by IH, B $\in K_\rho$, which, in turn, entails that B $\in$ C($K_\rho\cup$ {A}), and we are done.

Secondly we should check the revisions in K_{n+1}. Notice that for all E in K_{n+1}, C in L_0 and A in $L_>$, if C $\in$ s(E), then A $>$ C $\in K_s$. But then it is immediate, by construction, that C is in ρ(E) – because L_0 is included in $L_>$.

The constraint (c2) is trivially guaranteed by construction in the case of the revisions in K_{n+1}. With regard to expansions, take any E = $\langle Cn(K_\rho \cup \{A\}), C(K_s \cup \{A\})\rangle$. By IH $K_\rho \subseteq K_s$. But by the monotonicity of the underlying notion of consequence C: ρ(E) = Cn(K_ρ $\cup$ {A}) = C(K_ρ $\cup$ {A}) $\subseteq$ C(K_s $\cup$ {A}) = s(E). $\square$

References

[1] E. Adams, *The logic of conditionals* (Reidel, Dordrecht, 1975).

[2] H. Arló-Costa and I. Levi, Two notions of epistemic validity, *Synthese* **109** (1996) 217–262.

[3] H. Arló-Costa, Epistemic conditionals, snakes and stars, in: L. Fariñas del Cerro, G. Crocco, A. Herzig, eds., *Conditionals, from Philosophy to Computer Science* (Oxford University Press, Oxford, 1995) 193–239.

[4] Horacio Arló-Costa and Rohit Parikh, "On the inadequacy of of C2 (research note)", typescript, May 1995.

[5] H. Arló-Costa, Belief revision conditionals: models of suppositional reasoning, Report CMU-PHIL-84.

[6] C. Boutilier, 'Revision sequences and nested conditionals,' in: *Proceedings of International Joint Conference on Artificial Intelligence* (IJCAI, 1993).

[7] B. Chellas, Basic conditional logic, *Journal of Philosophical Logic* **4** (1975) 133 – 153.

[8] J. Collins, Belief Revision, Thesis, Princeton (1991).

[9] C.B. Cross, R. Thomason, 'Conditionals and knowledge-base update,' in *Belief Revision*, P. Gärdenfors (ed.), Cambridge Tracts in Theoretical Computer Science 29.

[10] A. Darwiche and J. Pearl, 'On the logic of iterated belief revision', *Artificial Intelligence*89 (1997) 1-29.

[11] M. Dummett, *The logical basis of metaphysics* (Duckworth, London, 1992).

[12] N. Friedman and J. Y. Halpern.: 1996 'Belief revision: a critique,' manuscript.

[13] P. Gärdenfors, *Knowledge in Flux* (MIT Press, Cambridge, Mass., 1988).

[14] P. Gärdenfors, S. Lindström, M. Morreau, W. Rabinowicz: 1989, 'The negative Ramsey test: Another triviality result', in *The Logic of Theory Change*, (eds.) Fuhrmann, A., Morreau, M., Springer-Verlag, pp. 129-134.

[15] A. Gibbard, Two recent theories of conditionals, in: W. Harper, R. Stalnaker, G. Pearce, eds., *Ifs* (Reidel, Dordrecht, 1980)

[16] S.O. Hansson, In defense of the Ramsey test, *Journal of Philosophy* **89** (1992) 522–540

[17] H. Katsuno and A. Mendelzon, On the difference between updating a knowledge base and revising it, in: P. Gärdenfors, ed., *Belief Revision* (Cambridge University Press, Cambridge, 1991) 183–203.

[18] K. Kelly, O. Schulte, V. Hendricks, 'Reliable belief revision', *Proceedings of the X International Joint Conference for Logic, Methodology and the Philosophy of Science*, Florence 1995.

[19] D. Lehmann, 1995, 'Belief revision, revised,' in *Proc. Fourteen International Joint Conference on Artificial Intelligence* (IJCAI'95), pages 1534-1540.

[20] S. Lidstrom, and W. Rabinowicz, 1995, 'The Ramsey test revisited,' in: L. Fariñas del Cerro, G. Crocco, A. Herzig, eds., *Conditionals, from Philosophy to Computer Science* (Oxford University Press, Oxford, 1995) 147–193.

[21] I. Levi, 1977, 'Subjunctives, dispositions and chances', *Synthese* vol. no. 34, pp. 423-455.

[22] I. Levi, 1988, 'Iteration of conditionals and the Ramsey test', *Synthese* vol. no. 76 49–81.

[23] I. Levi, 1991, *The fixation of belief and its undoing*, Cambridge University Press, Cambridge.

[24] I. Levi, 1996, *For the sake of the argument: Ramsey test conditionals, Inductive Inference, and Nonmonotonic reasoning* (Cambridge University Press, Cambridge, 1996).

[25] D. Lewis, 1976, 'Probabilities of conditionals and conditional probabilities', *Philosophical Review* vol. no. 85, pp. 297-315

[26] V. McGee, A counterexample to modus ponens, *Journal of Philosophy* **82** (1985) 462–471.

[27] C. Meek and C. Glymour, 1994, 'Conditioning and Intervening,' *Brit. J. Phil. Sci.*45, 1001 – 1021.

[28] F.P. Ramsey, *Philosophical Papers*, H. A. Mellor, ed., (Cambridge University Press, Cambridge, 1990).

[29] H. Rott, 1989, 'A nonmonotonic logic of belief revision,' in *The Logic of Theory Change*, (eds.) Fuhrmann, A., Morreau, M., Springer-Verlag, pp. 135-181.

[30] B. Skyrms, 1987, 'Updating, supposing and MAXENT,' *Theory and Decision*, 225 – 246.