Multidimensional Possible-world Semantics for Conditionals

Richard Bradley

London School of Economics and Political Science

Adams's Thesis is the hypothesis that the probabilities of indicative conditionals equal the conditional probabilities of their consequents given their antecedents.[1] The hypothesis is strongly supported by both introspection and by empirical evidence relating to the use of conditionals in hypothetical reasoning.[2] For this reason many philosophers have been attracted to the idea of using Adams's Thesis as a constraint on alternative truth-conditional accounts of the meaning of conditionals to that given by the material conditional, widely considered to be unsatisfactory. The first attempt to embed Adams's Thesis in a truth-conditional semantics was made by Robert Stalnaker in his 1968 and 1970 papers. And although a famous triviality result of David Lewis seemed to put paid to Stalnaker's own proposal, there have been subsequent attempts to do so by, amongst others, van

I am grateful to Frank Döring, Dorothy Edgington, Christian List and Olivier Roy, as well as an anonymous referee for Philosophical Review, for their helpful comments on earlier drafts of the paper.

[1] See Adams 1975 for the canonical statement and defence of this thesis. It was proposed prior to this by Jeffrey (1964).

Fraassen (1976), McGee (1989) and Kaufman (2005), all involving modifications of one kind or another to the possible-worlds framework used by Stalnaker.

Despite these attempts, many (and probably most) philosophers have taken the triviality results of Lewis and others to show that it is impossible to reconcile Adams's Thesis with any kind of truth-conditional semantics without making sentence meaning belief-dependent in an unsatisfactory way. I will argue to the contrary, proceeding as follows. The first two sections will describe the conflict between possible-world semantics and Adams's Thesis in more detail and survey some of the possible responses to it. The third will examine one rather sophisticated attempt to accommodate Adams's Thesis within a modified possible-worlds framework, one which treats conditionals as random variables taking semantic values in the unit interval. Although this attempt fails, the way it does so is instructive and serves as a spring-board for my own proposal, which is developed in the fourth and fifth sections. This proposal involves a somewhat different modification of possible-worlds models, namely one in which the semantic contents of sentences are represented by sets of vectors of possible worlds, rather than by sets of worlds.

In what follows we work with a background language L and a set W = $\{w_1, w_2, \ldots, w_n\}$ of possible worlds, assumed for simplicity to be finite (nothing of substance depends on this assumption). The power set of W - i.e. the set of all subsets of W - is denoted by Ω and the power set of any subset A of W, by Ω_A. By convention when p is a probability mass function on W, then P will be the corresponding

---

[2] The hypothesis has been pretty thoroughly investigated by psychologists working on conditionals. See for instance Over and Evans 2003.

probability function on Ω, such that the measure that P places on any set of worlds is the sum of the masses of its world elements, as measured by p.

Throughout I will use non-italic capitals to denote sets of possible worlds and italic capitals as sentence variables, reserving the symbols A, B and C for variables that range over factual sentences only (these being sentences in which the conditional operator introduced below does not occur). When the context makes for clear application of this convention, the set of worlds at which the factual sentence A is true will be denoted by the non-italic letter A, and vice versa. The symbols $\neg$, $\wedge$ and $\vee$ will respectively denote the sentential operations of negation, conjunction and disjunction. The symbol $\rightarrow$ will denote the sentential operation performed by the words "If ... then ..." in English conditional sentences. We will restrict attention in this paper to conditionals sentences with factual antecedent and consequent, i.e. what are typically termed simple conditionals.

The use of a single conditional operator may seem to prejudge an important question in the study of conditional sentences, namely whether or not the grammatical difference between indicative and subjunctive or counterfactual conditionals marks a fundamental semantic difference. Certainly there should be no denying that indicative and subjunctive versions of the same sentence can be evaluated quite differently, as is manifest in Adams's famous example of the difference between "If Oswald didn't kill Kennedy, then someone else did" (which is very probably true) and "If Oswald hadn't killed Kennedy, then someone else would have" (which is probably false).

A good theory of conditionals should be able to explain these differences. Equally it should be able to explain the many similarities in the behaviour of the two kinds of conditional. Some authors do so by postulating two different semantic operations with some common properties, others by postulating a single one with a parameter whose values will differ in the two cases. Although I will develop a theory of the second kind, the use of a single conditional operator is motivated by a concern to keep things simple, rather than to force such a unified treatment. Prejudgment can in any case be avoided by thinking of the arrow as a sentence-operator variable. In keeping with this, when discussion is restricted to indicative conditionals I will use the symbol '>' to denote the specific value taken by this operator in this case.

1. The Orthodoxy

To explain the difficulty in accommodating Adam's Thesis within truth-conditional semantics, let me start by sketching out the orthodox possible-worlds model of language meaning and use. For the purposes of this essay it is most usefully captured by four central propositions.

a) *Semantics*

The meanings of sentences are given by the conditions in which they are true, conditions being represented by possible worlds. More precisely, the semantic contents of the L-sentences can be specified by a mapping v from pairs of sentences and possible worlds to the set of permissible semantic values. If we let the semantic value assigned to sentence A at a world w be denoted by $v_w(A)$, then the core of the orthodoxy can be given by two propositions:

(1a) *Bivalence*: $v_w(A) \in \{0,1\}$

(1b) *Boolean Composition*:

$$v_w(A \wedge B) = v_w(A) \times v_w(B)$$
$$v_w(\neg A) = 1 - v_w(A)$$
$$v_w(A \vee B) = v_w(A) + v_w(B) - v_w(A \wedge B)$$

Bivalence says that sentences can take only one of two possible semantic values – truth (1) or falsity (0) – at each possible world. The meaning of the sentence A can therefore be identified with the set of worlds in which it is true, denoted hereafter by $[A]$. Boolean Composition, on the other hand, determines the relation between the semantic values (truth-conditions) of compound sentences and those of their constituents.

b) *Pragmatics*

The degree to which a rational agent will believe a sentence is given by her subjective probability for the sentence being true. More formally, let p be a probability mass function on the set of worlds that measures the probability of each world being the actual world. Then the rational agent's degrees of belief in sentences will equal her expectation of their semantic value, $\mathrm{E}(v(A))$, i.e. be given by a probability function Pr on L such that for all L-sentences A:

$$(2)\ \Pr(A) = \mathrm{E}(v(A)) = \sum_{w \in W} v_w(A).p(w)$$

In virtue of Bivalence this implies that:

$$\Pr(A) = \sum_{w \in [A]} p(w) = P([A])$$

c) *Logic*

A sentence B is a semantic consequence of another sentence A (denoted $A \models B$), relative to the value assignment v, just in case the truth of A ensures that of B. Formally:

(3) $A \models B$ iff $[A] \subseteq [B]$

Note that (1a) and (1b) together with (3) ensures that $\models$ is a classical consequence relation.[3]

d) Explanation

The final claim concerns the relationship between the semantics, pragmatics and logic of a language. Loosely, it is this: what belief attitudes it is rational to take to sentences and what inferences it is correct to make with them is determined by what sentences mean and what beliefs one has about possible worlds, and not the other way round.

To make this more precise, let $\Pi = \{p_j\}$ be the set of all probability mass functions on the set of possible worlds W, interpreted as the set of rationally permissible beliefs. And let $V_L = \{v_i\}$ be the set of all permissible assignments of semantic values to sentences of L.[4] A possible-worlds model (PW-model for short) of L is a structure <W, v, p> where W is the background set of worlds, v belongs to V_L and p to Π. Such a structure determines both what belief attitudes the speaker can rationally take to L-sentences and what inferences she can rationally make with

[3] It is (1b) Boolean compositionality that is doing the work here. Bivalence is required only because of the specific manner in which logical consequence is defined in (3).

[4] What assignments are permissible depends on the semantic theory under consideration and in particular what values the theory requires of compounded sentences. Condition (1b) above constrains V_L to contain only those assignments respecting the Boolean laws, but does not constrain the assignment to conditionals in any other way.

them. In particular, if *Pr* and $\vdash$ are respectively a probability measure and a consequence relation on L-sentences then we can say that a PW-model <W, *v*, *p*> explains the pair (Pr, $\vdash$) just in case Pr and $\vdash$ are related to *v* and *p* by (2) and (3). That is, it yields explanations of the form '$A \vdash B$ *because* $[A] \subseteq [B]$' and '$Pr(X) = \mathrm{x}$ *because* $P([\mathrm{X}]) = \mathrm{x}$'.

The final assumption underlying standard applications of possible-worlds models can now be made explicit:

(4) For all $v \in V_L$ and $p \in \Pi$, $< W, v, p >$ is a PW-model of L.

The implication of (4) is two-fold. Firstly, the semantic assignment is independent of the agent's belief over worlds, and vice versa. And secondly, there are no constraints on agents' attitudes to sentences other than those contained in the specification of V_L and Π. At first sight the latter claim might seem to be implausibly strong because it is too permissive about what counts as rational belief. For instance, strong believing that one knows that it is soup for dinner should preclude only weakly believing that it is soup for dinner. Similarly, believing that the chance of a coin landing heads is 0.9 precludes only weakly believing that the coin will land heads. The thought behind (4) is that such constraints on rational belief should be implied by the relation between the semantic values assigned to sentences expressing claims about, say, knowledge or chances and those assigned to the sentences picking out the events which are subject to chance or are the object of knowledge. So these examples don't count against it. Nonetheless, I will argue later on that there are other reasons for weakening (4).

Adoption of the possible-worlds framework just described is of course consistent with numerous different specific theories about the truth conditions of conditionals. The focus of our interest is on those that are compatible with a highly plausible constraint on rational belief known as the Ramsey Test hypothesis. The Ramsey Test hypothesis says that your belief (or degree of belief) in a conditional $A \rightarrow B$ should match or equal your belief (or degree of belief) in the consequent B, on the supposition that the antecedent A is true. To apply it one must be able to determine what one would believe under a supposition. To make things more complicated, there is more than one way of supposing that something is true. One way is to suppose that, as a matter of fact, something is true, such as when we assume that Oswald was not in fact the one who killed Kennedy. This will be called matter-of-fact or *evidential* supposition. Alternatively one might suppose that something is true, contrary to what one knows or believes to be the case, such as when we suppose that, for the sake of argument, Oswald hadn't killed Kennedy. This will be called contrary-to-fact or *counterfactual* supposition. (Possibly there are more kinds of supposition, but all that matters to our discussion is that there are at least two distinct kinds).

One of the great strengths of the Ramsey Test hypothesis is that these different ways of supposing something true match the differences in our evaluations of indicative and subjunctive conditionals: roughly, we determine our beliefs in indicative conditionals by evidential supposition of the antecedent, and of subjunctives by counterfactual supposition. Furthermore, it is generally accepted that when belief comes in degrees, evidential supposition is achieved by ordinary Bayesian conditioning on the sentence that is supposed true. That is, when I suppose that as a matter of fact that X, I adopt as my (suppositional) degrees of belief my (pre-suppositional) conditional degrees of belief, given that X. It follows that if my degrees of belief in factual sentences are given by probability measure Pr, then my degree of belief in the simple indicative conditional $A>B$ should be given by:

Adams's Thesis: $\Pr(A>B) = \Pr(B|A)$

A corresponding application of the Ramsey Test hypothesis for degrees of belief to counterfactual conditionals would yield that they should be believed to a degree equal to the probability of truth of their consequent on the supposition that, contrary-to-fact, their antecedent were true.[5] Proposals as to how such supposition works do exist, including those of Adams (1975), Lewis (1976) and Skyrms (1981, 1994), and although our focus will be on indicative conditionals, nothing about the semantic model that will be developed here will rule them out. In particular, what we might call Skyrms's Thesis, namely the claim that the probability that of 'If *A* were the case, then *B* would be' equals the expected chance of B given that A, will turn out to be a consequence of the model under some natural assumptions about counterfactual supposition.

Adams's Thesis is widely recognised both to capture our intuitions about rational belief in conditionals and to provide the best explanation for the empirical evidence concerning the role played by conditionals in the inferences that people make. But, as mentioned before, a series of triviality results show that it is impossible to accommodate Adams's Thesis within the kind of possible-worlds semantic

[5] Not everyone agrees that the Ramsey Test hypothesis should apply to counterfactuals. Lewis's truth conditions for counterfactuals, for instance, imply that many highly probable conditionals (according to the Ramsey Test) are false. I regard this as a weakness of Lewis's theory however. See Edgington (1995) for some convincing counterexamples to Lewis's theory.

framework described above.[6] Indeed it is not even possible to accommodate a very weak consequence of Adams's Thesis, known as the Preservation condition, without generating highly implausible conclusions. This latter condition says that if it is epistemically impossible that *B*, but possible that *A*, then it is epistemically impossible that if *A* then *B*. Formally:

Preservation condition: If $\Pr(A) > 0$, but $\Pr(B) = 0$, then $\Pr(A>B) = 0$.

Let us examine more precisely why the Preservation condition, and hence Adams's Thesis, is not non-trivially consistent with the orthodoxy as presented here. Let *A* and *B* be any factual sentences in L. Then the Preservation condition implies in conjunction with (1a) and (2) that for all PW-models of L, $\langle W, v, p \rangle$, such that $P([B]) = 0$ and $P([A]) > 0$, it must be the case that $P([A>B]) = 0$. And in virtue of (4) this can be so only if for every semantic assignment v, either $[A] \subseteq [B]$ or $[A>B] \subseteq [B]$. If $[A] \subseteq [B]$, then the antecedent of the Preservation condition is never satisfied, so no constraint is imposed on the probability of $A>B$. If $[A>B] \subseteq [B]$ then the laws of probability require that $\Pr(A>B) = 0$ whenever $\Pr(B) = 0$. But if neither of these conditions hold, then (4) ensures the existence of a probability p on worlds assigning zero weight to $[B]$, but non-zero weight to both $[\neg B \wedge A]$ and $[\neg B \wedge (A>B)]$ – in violation of the Preservation condition. It then follows by (3) that the Preservation condition can be satisfied by all PW-models of L only if L is trivial in the sense of containing no sentences *A* and *B* such that *B* is logically independent of both *A* and $A>B$.[7]

---

[6] In addition to Lewis (1976), see, for instance, Hajek (1989), Edgington (1991), Döring (1994), Bradley (1998) and the papers by Hajek and Hall in Eels and Skyrms (1995).

[7] See Bradley (2000) for a more formal statement of this argument.

This argument makes essential use of claims (1a), (2) and (4). Claims (1a) and (2) together ensures that the probabilities of sentences equal the probabilities of the sets of worlds in which they are true, and (4) ensures that the probabilities of these sets of worlds are independent of the assignment of them to sentences. It follows that if the Preservation condition is true then either conditionals are not just true or false at a world, or that their probabilities are not probabilities of truth, or that there is some restriction on the co-assignment of meaning to sentences and beliefs to worlds not contained in the standard theory.

2. Routes to Reconciliation

A wide range of responses to this problem has been explored in the literature. Authors such as Lewis (1976), Jackson (1979) and Douven (2007) argue that the triviality results show that Adams's Thesis is false as a claim about rational belief and that the evidence we are disposed to assert conditionals to a degree equal to the conditional probability of their consequent given their antecedent should be explained by pragmatic principles of one kind or another, not by the semantic content of indicative conditionals (which they take to be that of the material conditional). But as I have argued elsewhere (Bradley 2002) these accounts are rather unsatisfactory because that don't extend in a natural way to sentences containing conditionals. For example the sentence "If I try to climb Mt. Everest, then I will succeed" is, on these accounts, very probably true, (because I won't attempt the ascent) but not necessarily assertable. But then why is the sentence "It is probable that if I try to climb Mt. Everest, then I will succeed" also not assertable?[8]

At the other extreme, and a good deal more plausibly, non-factualists such as Edgington (1991, 1995) and Gibbard (1981) argue that the triviality results show that conditionals don't make factual claims and hence do not have (standard) truth conditions. This response implies giving up the possible-worlds framework entirely and adopting Adams's Thesis as a stand-alone hypothesis about rational belief in conditionals. The problem with strategy is that it makes it something of a mystery that we argue over the claims expressed by conditional sentences in much the same way as we argue over factual claims (i.e. by arguing over what is the case, not over what we believe to be the case). Furthermore, without some account of semantic value, it is difficult to explain how we compound conditionals with other sentences using the usual sentential connectives and how we can make inferences with conditionals that eventuate in sentences that make factual claims.[9] (Consider Modus Ponens: how can we infer the truth of B from that of A using the hypothesis that $A>B$, if the latter makes no truth claim?)

In an earlier paper (Bradley 2002), I argued for an intermediate position, namely that conditionals do take truth values, but only in worlds in which their antecedents are true. Others (e.g. Milne 1997, McDermott 1996) have pursued a similar strategy, dropping Bivalence in favour of a three-valued semantics based on

---

8 The obvious answer is because it is false, as Adams's Thesis would imply.

9 Neither Gibbard nor Edgington are impressed by this argument as they doubt that any truth-functional theory does a good job of explaining the way in which we actually compound conditional sentences.

the values of truth, falsity and neither. This approach is able to deliver an explanation of Adams's Thesis by modifying (2) in favour of the hypothesis that probabilities of sentences are their probabilities of truth, conditional on them being candidates for truth i.e. either true or false. Unfortunately, in my opinion, nobody pursuing this strategy has given a convincing account of how the truth-values, and hence probabilities, of compounded sentences depend on those of their constituents.

In this paper, I want to look at another class of responses, namely those that involve, in one way or another, some restriction on the co-assignment of meaning to sentences and beliefs to worlds. There are two salient candidates for restrictions. It might be the case that what sentences mean depends on what beliefs one holds, or vice versa. This possibility is explored in the papers of van Fraassen (1974), McGee (1989), Jeffrey (1991), and Jeffrey and Stalnaker (1994) for instance. Alternatively it might be the case that there are restrictions on what beliefs we can hold that are not contained in the requirement that degrees of belief be probabilities. Either way (4) would fail: some probability functions on worlds would not be admissible belief measures or some combinations of meaning and belief would be impossible.

Let us start by looking at a theory of the first kind, that of Jeffrey and Stalnaker (1994), in which the contents of sentences are treated, not as bivalent propositions, but as random variables taking values in the interval [0,1]. In particular, the content of a simple conditional $A \rightarrow B$ is represented as a random variable taking the value '1' in all worlds in which both A and B are true, the value '0' in all worlds in which A is true, but B is false, and the conditional probability of B given A in all worlds in which A is false. More formally, a semantic assignment is for them a mapping from sentences to [0,1] satisfying Boolean Compositionality plus:

JS-semantics:

$$v_w(A \to B) = \begin{bmatrix} v_w(B) & \text{if } v_w(A)=1 \\ \mathrm{E}(v(B \mid v(A)=1)) & \text{if } v_w(A)=0 \end{bmatrix}$$

The subjective probabilities of sentences are still determined by their semantic values in the 'orthodox' way – i.e. in accordance with equation (2) - with the probability of a sentence being its expected semantic value. It then follows from the JS-semantics that:

$$\Pr(A \to B) = \Pr(A \wedge B).1 + \Pr(A \wedge \neg B).0 + \Pr(\neg A).\Pr(B \mid A) = \Pr(B \mid A)$$

in accordance with Adams's Thesis. It also follows that a rational agent's beliefs will satisfy a condition that will prove important later on:

Independence: If $C \vdash A$ then $\Pr(C \wedge (A \to B)) = \Pr(C).\Pr(A \to B)$

The most notable feature of Jeffrey and Stalnaker's account is that the meanings of conditionals depend on agents' subjective degrees of belief, since the latter determine the semantic value of a conditional in worlds in which its antecedent is false. Hence the contents of conditionals are not strictly random variables (despite the title of their paper), but functions from probability measures to random variables.[10] Because belief restricts meaning, condition (4) is violated, and it is this feature of their account, rather than the dropping of Bivalence, that enables them to satisfy the Preservation condition without running afoul of our triviality result. For, on their account, the element in the Boolean algebra of random variables that is picked out by the sentence $A \to B$ varies with the agent's beliefs. In particular when $\Pr(B) = 0$, the content of $A \to B$ is the same random variable as that of $A \wedge B$, so that

[10] This is somewhat obscured by the fact that Jeffrey and Stalnaker implicitly assume a fixed probability measure in their discussion.

every probability measure on the algebra of random variables determined by Pr must give measure zero to the random variable associated with the sentence $A \rightarrow B$. Hence the Preservation condition is non-trivially satisfied on their account.

Jeffrey and Stalnaker's account does leave some questions unanswered however. Firstly, since random variables take values other than '0' and '1' (or 'false' and 'true') probabilities of sentences are not ordinary probabilities of truth. How then are these probabilities, and hence Adams's Thesis, is to be interpreted? Secondly, and more importantly, they do not offer any explanation as to why the semantic values of conditionals should be related to the agent's conditional beliefs in the way that they postulate and how this dependence of meaning on belief is to be squared with the fact that we appear to use conditionals to make claims about the way that the world is rather than to express our state of mind.

Partial answers to both these questions can be drawn, I believe, from an earlier paper of Vann McGee (McGee 1991). McGee adopts a modified version of Stalnaker's (1968) semantic theory for conditionals in which a conditional $A \rightarrow B$ is true at a world w iff its consequent, B, is true at the world f(w, A), where f is a selection function picking out for any sentence and world w, the 'nearest' world to w at which the sentence is true.[11] Which world is the nearest is not something that is determined by the facts, but depends in part on the agents' beliefs. According to McGee (1989, 518),

> Purely semantic considerations are … only able to tell us which world is the actual world. Beyond this, to try to say which of the many selection functions that originate at the actual world is the

[11] McGee's modification will not matter here, so for simplicity I will omit it.

actual selection function, we rely on pragmatic considerations in the form of personal probabilities.

In general these pragmatic considerations will be insufficient to determine which world is the nearest one in which the antecedent of a conditional is true (or which selection function is the right one). Hence they will not determine the truth-value of a conditional at a world. But the agent's partial beliefs will constrain the choice of selection function to the extent of determining what might be called the expected truth-value of a sentence at a world, where the latter is defined as the probability weighted sum of the truth-values ('0' or '1') of a sentence, given particular selection functions, with the weights being determined by the subjective probabilities of the selection functions. More formally, let $F=\{f_i\}$ be the set of selection functions on $W\times\Omega$, with $f_i(w, A)$ being the world selected by f_i as the nearest one in which A is true. Let q be a probability mass function on F. Then the semantic value of a conditional at a world w is given by:

$$\text{(1c)} \qquad v_w(A \to B) = \sum_{f_i} v_{f_i(w,A)}(B).q(f_i)$$

The expected truth-values yielded by agents' uncertainty about distance between worlds provide a natural interpretation of the intermediate semantic values postulated by Jeffrey and Stalnaker. However equation (1c) does not by itself imply that conditionals take the specific semantic values claimed by JS-semantics. For this the probabilities of selection functions codifying possible nearness judgments must be correlated with the agent's degrees of conditional belief in the right kind of way. On McGee's account this correlation is secured by means of further constraints on rational partial belief, most notably including the aforementioned Independence principle. The additional constraints on rational belief suffice to determine probabilities for selection functions, identified by McGee with complex conditionals of the form $(A\rightarrow w_A) \wedge (B\rightarrow w_B) \wedge (C\rightarrow w_C) \wedge \ldots$, i.e. explicit descriptions of the world that would be the case for each possible condition supposed true. Equations (1c) and (2) then yield Adams's Thesis.

To summarise: on this interpretation of McGee's theory, the semantic value of a conditional at a world at which its antecedent is false is its expected truth value, calculated relative to probabilities of selection functions. Probabilities of selection functions are in turn determined by the agent's partial beliefs, assumed to conform to both the laws of probability and the Independence principle. Hence it is these latter properties of rational belief that determine what conditionals mean and explain why rational belief in conditionals must satisfy Adams's Thesis.

3. Testing McGee's theory

The Independence principle is unsatisfactory in one notable respect (at least from the perspective of the orthodoxy): it is stated as a constraint on belief attitudes towards sentences rather than their contents. Ideally, however, our attitudes to sentences should be explained in terms of what these sentences mean or say and what belief attitude it is rational to take to such meanings. Since, on McGee's account, the semantic value of a conditional depends not only on what world is the actual one, but also on what world is picked out by selection functions, it would be expected that such an explanation would make reference to the properties of rational belief attitudes to selection functions. Instead McGee simply argues that the Independence principle has to be true if Adams's Thesis is. For were it not, Adams's Thesis would not survive belief change by Jeffrey conditionalisation.

This argument is not persuasive. As Richard Jeffrey himself pointed out, rationality does not require belief revision by Jeffrey conditionalisation. It is only rationally mandatory when the shift in probability of some proposition X, does not disturb the conditional probabilities of any other proposition conditional on X. But this is precisely what must be assumed for McGee's argument to go through.[12] However, rather than offering further theoretical considerations for or against the Independence principle, I propose to test both it and Adams's Thesis against a couple of examples in which application of the central concepts of McGee's semantic theory – possible worlds and selections functions – is unproblematic.

The first example is a case in which there is uncertainty about what the correct selection function is, but not uncertainty about which world is the actual one. Suppose that we have before us a coin that is known to be fair, so that the chance of it landing heads in the event of being tossed is 0.5. This set-up can be illustrated as follows.

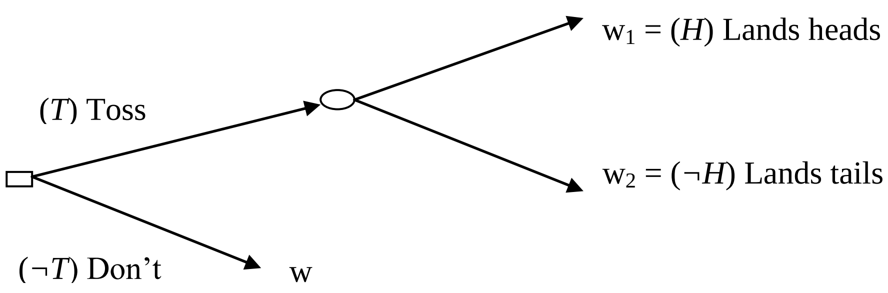

Example 1: Selection Uncertainty

[12] An associated argument for a more complex version of the Independence principle is made in terms of fair betting arrangements, but this argument assumes the simpler version of the principle.

What is the probability that if the coin is tossed it will land heads? There are only two plausible selection functions to consider, namely functions f and f′ such that:

f(w1,T) = w1, f(w2,T) = w2, f(w3,T) = w1

f′(w1,T) = w1, f′(w2,T) = w2, f′(w3,T) = w2

If we calculate the probability of $T \rightarrow H$ using (1c′), we see that Adams's Thesis is satisfied, irrespective of the probability of the coin being tossed, since:

$p(T \rightarrow H) = p(\text{w1}) + p(\text{w3}).q(\text{f}) = 0.5$

$p(H \mid T) = p(\text{w1})/(p(\text{w1}) + p(\text{w2})) = 0.5$

In this simple case of selection uncertainty, therefore, his theory offers a plausible explanation for the truth of Adams's Thesis.

The second example is a case in which there is uncertainty about which world is the actual one, but no uncertainty about what the correct selection function is.[13] Suppose that we have before us a coin that is known to be biased, but that we do not know whether it is biased in favour of heads or in favour of tails (it is either a two-headed or two-tailed coin say).

[13] This second example has much the same structure as the (more colorful) one used by Mark Lance (1990) and leads to much the same conclusion.

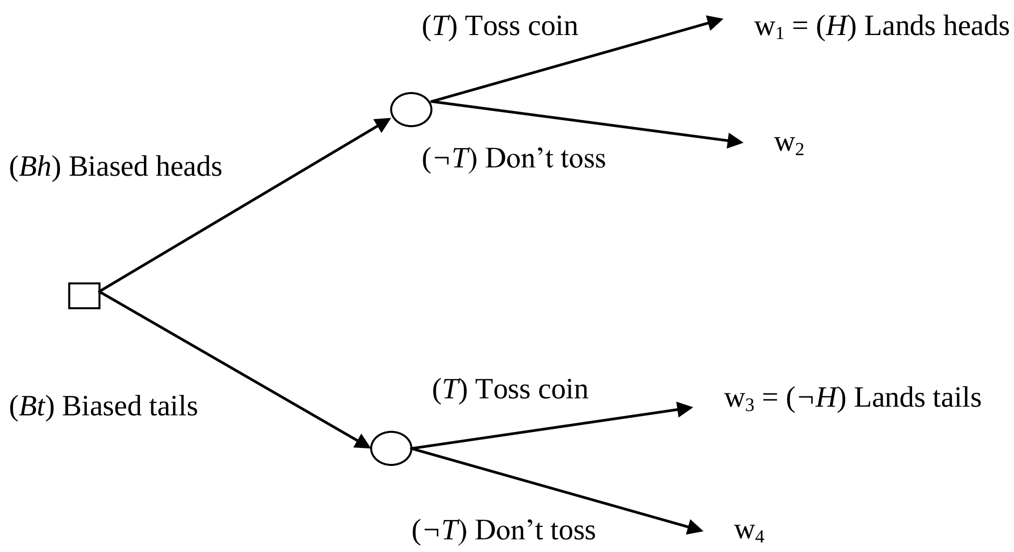

Example 2: World Uncertainty

In worlds in which the coin is tossed, if it is biased heads it will land heads, and if biased tails it will land tails. What about worlds in which it is not tossed? In view of the known fact of bias, there is only one plausible selection function in this case: it is the function f such that:

f(w2,T) = f(w1,T) = w1

f(w4,T) = f(w3,T) = w3

It follows that:

$$P(T \to H) = p(w_1) + p(w_2) = P(Bh)$$

Likewise, on the assumption that whether the coin is biased one way or another is independent of whether it is tossed (which is a reasonable assumption to make in this context):

$$P(H \mid T) = \frac{p(w_1)}{p(w_1) + p(w_3)} = P(Bh \mid T) = P(Bh)$$

So in this example too Adams's Thesis follows by application of (1c).

On the other hand, the Independence principle fails. To see this note that Independence requires that at every world in which the coin is not tossed, the probability that it would have landed heads had it been tossed, is probabilistically independent of it being tossed. So by Independence:

$$\Pr(T \to H \mid \neg T \wedge Bt) = \Pr(T \to H)$$

But:

$$\Pr(T \to H \mid \neg T \wedge Bt) = 0$$

because, given that the coin is biased towards tails, it is certain that the coin would not have landed heads had it been tossed. On the other hand, as we have seen:

$$\Pr(T \to H) = \Pr(Bh) > 0$$

But this contradicts the above. So Independence must be false.

The Independence principle is closely related to a claim common to McGee and Jeffrey and Stalnaker's theories, namely that the semantic value of a conditional equals the conditional probability of its consequent given its antecedent in all worlds in which its antecedent is false. This claim too is not satisfied in the second example: the semantic value of $T \to H$ at world w4 is not 0.5 but 0, since at this world it is certain that at the nearest world in which T is true (namely w3), it is false that H. It would seem that the JS-semantics upon which these theories rest is unsatisfactory as a claim about the contents of indicative conditionals.

Fortunately for the project of accommodating Adams's Thesis it is not necessary that the semantic value of a conditional be the conditional probability of its consequent given its antecedent at *every* world in which its antecedent is false. All that is required is that *on average* it has this semantic value. Or to put is slightly differently, what is required for Adams's Thesis is not Independence, but the following, logically weaker, condition:

Restricted Independence: $\Pr(\neg A \wedge (A \rightarrow B)) = \Pr(\neg A).\Pr(A \rightarrow B)$

And this condition, like Adams's Thesis (but unlike Independence) does not seem to be undermined by our examples. So in the next section, I would like to build on what we have learnt from these theories to construct a more plausible semantic basis for this restricted kind of independence.

4. Two-dimensional Semantics

The discussion thus far suggests that two types of uncertainty are at play when evaluating conditionals. On the one hand there is the familiar uncertainty about what is the case or about which world is the actual one. On the other hand, there is uncertainty about what is or would be the case if some supposed condition is or were true, or about which world is the nearest one satisfying the condition. I will speak in this latter case of uncertainty as to which world is the counter-actual world under the supposition in question.

McGee represents the first kind of uncertainty in the standard way, by a probability mass function on worlds measuring the probabilities of them being the actual world. The second kind is measured by a mass function on selections functions, each of which represents a hypothesis as to which world is the counter-actual one under each possible supposition. The mass on a selection function gives the probability that it correctly identifies these counter-actual worlds.

To see how this works, consider the following (extremely simple) possible-worlds model shown in the diagram below, exhibiting a set of four possible worlds

W = {w1, w2, w3, w4} and 15 associated non-empty propositions, with for instance A={w1, w2} and B={w1, w3}.

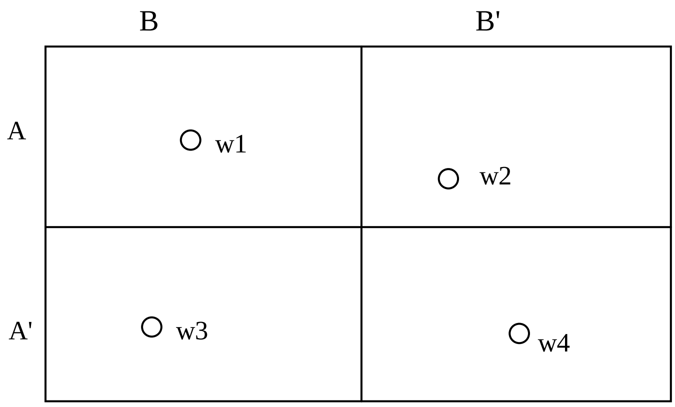

Even such a simple model has quite a few possible selection functions associated with it, since each world and proposition pair can take any of the four worlds as values (subject to the usual constraints on selection functions). To derive the probability of the conditional $A \rightarrow B$ in McGee's framework, for instance, we need to determine its expected truth value at each world and then calculate the average of its expected truth values, using the probabilities of worlds as weights. But to work out the expected truth value of $A \rightarrow B$ at w3 say, we need to assign probabilities to possible function f: f(w3, A) = w1, f(w4, A) = w1; function g: g(w3, A) = w1, g(w4, A) = w2; and so on. And then use these as weights in finding the average truth value of the conditional at that world.

This is unnecessarily complicated. Instead of trying to judge the probability of every complete specification of distance between worlds, we can turn things around and simply judge the probability that each world is the nearest one to the actual world in which the antecedent of the conditional is true. When evaluating the conditional $A \rightarrow B$, for instance, we can confine attention to the two worlds w1 and w2 in which *A* is true and ask ourselves how probable it is that each is or would be the counter-actual world, if *A* is or were true. These probabilities may then serve as the weights on the truth-value of $A \rightarrow B$ at w1 and w2 (1 and 0 respectively) that are needed to calculate the expected truth value of the conditional.

Note that how we judge this kind of uncertainty will depend on the kind of supposition that we are engaging it. When evaluating an indicative conditional by engaging in evidential supposition, we need to evaluate how probable it is that one or another A-world (really) is the case, given the truth of *A*. On the other hand when we are engaged in the kind of contrary-to-fact supposition appropriate to the evaluation of counterfactuals, we should evaluate how probable it is that each world would be the case were *A* true (contrary-to-fact).

To summarise: what I am suggesting is that we represent the second kind of uncertainty by a probability function on *worlds*, not on selection functions. This function on worlds measures not their probability of being the actual world but their probability of being the counter-actual world under a supposition. This has the advantage of allowing us to dispense altogether with talk of selection functions and simply work with possible worlds.

With this simplification in place, we can represent in the following way our state of uncertainty (or the part of it relevant to the evaluation of the conditional $A \rightarrow B$) for the simple model we have been using. There are four possible worlds that we need to assess with regard to their probability of being the actual world, and two possible A-worlds which need to be assessed with regard to the probability that they are or would be case if *A* is or were (worlds at which *A* is false are not candidates for this). What we need to examine is how these assessments depend on each other. For this purpose we can tabulate their objects in the following way.

	Supposed A-Worlds	
Possible Worlds	**W1**	**W2**
W1	<w1,w1>	<w1,w2>
W2	<w2,w1>	<w2,w2>
W3	<w3,w1>	<w3,w2>
W4	<w4,w1>	<w4,w2>

Each ordered pair $<w_i,w_j>$ in the table represents a possibility: the event that w_i is the actual world and that w_j is the counter-actual A-world. Sets of such possibilities will serve as the meanings or contents of sentences. The contents of *factual* sentences are given by rows of the table. The sentence *A*, for instance, has as its content the first and second rows of the table, while sentence *B* has the first and third. The contents of *conditional* sentences, on the other hand, are given by columns of the table. The sentence $A \rightarrow B$, for instance, has as its content the first column of the table and that of $A \rightarrow \neg B$ the second. The contents of conjunctions, disjunctions and negations of sentences (conditional or otherwise) are given by the intersection, union and complements of the contents of their component sentences.

When a pair of worlds $<w_i,w_j>$ is part of the content of a sentence *X* we can say that $<w_i,w_j>$ makes *X* true as a kind of short-hand for the claim that w_i being the actual world and w_j being the counter-actual A-world makes it true that *X*. In this sense the truth conditions of sentences are given in this model by ordered pairs of worlds. An immediate implication is that we cannot in general speak of a sentence as being true or false at a world *simpliciter*. For instance, while $A \rightarrow B$ is true at <w3,w1>, we cannot say that it is made true (or false) by the facts at w3 because these facts alone do not determine its truth or falsity independently of the relevant counter-facts i.e. of whether the counter-actual A-world is w1 or w2.

The upshot is that the theory being proposed gives a truth-conditional semantics for conditional sentences, while at the same time allowing that the truth-values of conditionals are not determined by the facts alone. The former property allows for an explanation of the role that conditionals play in discourse aimed at establishing the truth and of how they compound with other sentences. The latter property explains the difficulty we have in some situations in saying whether a conditional is true or false, a difficulty that Lewis used to motivate his rejection of the law of conditional excluded middle and that non-factualists point to in motivating their rejection of truth-conditional accounts of conditionals.[14] It is also what makes the accommodation of Adams's Thesis possible, as we shall see in the next section.

To insist on the formal distinction between facts and counter-facts is not to deny that they might be related in various specific ways, both semantically and

---

[14] Consider Gibbard's (1981) story in which Jack and Zack watch Sly Pete play poker. Before the room is cleared Zack sees that Pete can see his opponent's hand and declares later that if Pete called he won. Jack sees that Pete has the losing hand and declares later that if Pete called he lost. Suppose Pete didn't call. Who was right: Jack or Zack? Like Gibbard I would claim that this isn't settled by what we have been told, since the reported facts don't suffice to determine the counter-facts under the supposition that Pete called. But unlike Gibbard, I take this to be consistent with the two declarations being either true or false (relative to the counter-facts).

pragmatically. Any view of their relation must square with the fact there are questions such as whether the allies would have lost the second world war if Hitler had captured Moscow, or whether I would have been a philosopher if I had been born in a different family, that seem impossible to settle no matter how much evidence we can bring to bear on them concerning what actually happened. On the other hand there are questions such as whether the sugar would have dissolved if I had added it to my coffee that do seem to be decided by features of the actual world: the chemical properties of the sugar, the temperature of the coffee, how much sugar had already been added, and so on.

There are two extreme views that fare badly in this regard because they have trouble explaining one of these two classes of cases. On the Autonomy view the counter-facts are completely independent of the facts – they are 'barely' true, to borrow a phrase of Dummett - and hence any combination of facts and counter-facts is possible. On the Reductionist view, on the other hand, the counter-facts are completely determined by the facts. The Reductionist view leads to a dead-end as far as the project of this paper is concerned, because it takes us back to the orthodoxy and incompatibility with Adams's Thesis. Not so, the Autonomy view. But it has an implausible implication, namely that whether the sugar will dissolve if added to the coffee is independent of whether or not the sugar was in fact added to the coffee and whether or not it dissolved when added.

More promising is some kind of intermediate non-reductionist view that recognises a variety of possible relations between facts and counter-facts. Some classes of counter-facts might be completely determined by a broad enough specification of the facts pertaining at a possible world (as in the sugar case). Others might be constrained, though not fully determined, by the chances of relevant events. For instance, in the second coin tossing example, the fact that a coin was biased heads was taken as grounds for saying that it would have landed heads if it had been tossed. Still others might be hardly constrained at all. Views about these relations can be accommodated in two ways: by restricting the possible combinations of facts and counter-facts serving as semantic values of sentences and by restricting the joint attitudes that an agent can rationally take to them. Discussion of the latter class of restriction must be deferred until the next section. For now we focus on a widely adopted semantic principle - that I will call Centring in line with the terminology introduced by David Lewis – and which suffices to rule out the Autonomy view.

According to Centring, if world w_i is the actual world and *A* is true at w_i, then w_i is also the counter-actual A-world. For instance if I added sugar to my coffee and it dissolved then, under the supposition that I added it to my coffee, the sugar dissolved (after all I *did* add it). In our toy model the effect of adopting Centring is to eliminate some cells from the table of possibilities, leaving:

	Supposed A-Worlds	
Possible Worlds	**W1**	**W2**
W1	<w1,w1>	-
W2	-	<w2,w2>
W3	<w3,w1>	<w3,w2>

W4	<w4,w1>	<w4,w2>

Centring also has important implications for the logic of conditionals implied by the proposed construal of semantic content. To obtain it we define semantic entailment in the same way as before, via claim (3) of the orthodoxy (namely that $A \models B$ iff $[A] \subseteq [B]$), except that the function [.] assigns not sets of worlds, but sets of pairs of worlds (more generally sets of world vectors) to sentences. Then it follows from way in which contents are assigned to sentences that the consequence relation $\models$ so defined will validate the laws of classical logic as well as a number of principles of conditional logic: $A \to B \models A \to (A \wedge B)$, $B \models A \to A$ and so on. If we assume Centring then several further properties are satisfied, most notably that $A \wedge B \models A \to B \models \neg A \vee B$ and that, in accordance with Modus Ponens, $A \wedge (A \to B) \models B$.[15]

It is noteworthy that on any semantic assignment, irrespective of whether Centring holds or not, it will be the case that $[\neg(A \to B)] = [A \to \neg B]$. Since the law of conditional excluded middle which it implies has been the subject of much debate, it is worth reflecting on why this identity holds. The elementary possibilities that we are working with are maximally specific descriptions of both how things are and how they would be if it were the case that *A*. If it were the case that *A* then things would be one way or another: the counter-facts would be such that *B* was the case or they would be such that *B* was not. Lewis (1973) offers apparent counter-examples to this claim. He argues, for instance, that neither would both Verdi and Bizet have been French if they had been compatriots, nor would they both have been Italian. But while it is entirely plausible that the facts in this case (including the similarities that Lewis appeals to) do not determine the relevant counter-facts of co-nationality, the latter must go one way or another: either Verdi and Bizet would have been compatriots by being both French or by being both Italian (or, of course, by both being something else).

The point is relevant to some of the broader philosophical questions that the proposed model gives rise to. The central feature of the formal semantics is the distinction between facts and counter-facts for it is this that allows the ascription of non-factual truth conditions to conditionals. It is important to note that neither this distinction, nor the associated claim that the counter-facts are not reducible to the facts, implies any position on the metaphysical and epistemological status of the counter-facts. In particular, in making the distinction I am not taking a stand on whether or not counter-facts are real in the sense of being objective features of the world that we inhabit. Nor, in insisting on the non-reducibility of the counter-facts to the facts, do I intend to commit myself to a view as to whether or to what extent we can have knowledge of the counter-facts. Indeed questions about what is real and/or knowable seem to me to cut across the distinction between facts and counter-facts

---

[15] The fact that Modus Ponens depends on Centring is important for the generalisation of the model to nested conditionals, since it offers a possible explanation for kind of failure of Modus Ponens for inferences with right-nested conditionals that McGee identified.

since we can as well ask whether certain kinds of facts (for instance, those concerning chances or the future) are objective and determinable.[16]

The advantage of working with a formal semantics that is largely neutral on these kinds of questions is that it can be adopted by those holding quite different philosophical views, thereby providing a common framework within which their respective positions can be debated. Realists can construe both the facts and counter-facts that appear in it as different features of reality about which we can be uncertain but which can be investigated. Anti-realists can construe the use of counter-facts to fix the truth-values of conditionals as merely a formal device to support a compositional semantics. Both can use it to explore the relation between the content of conditionals, the attitudes we take to them and kinds of uses to which they are put.

But does the way in which truth conditions are ascribed to conditionals not entail realism about the counter-facts? Not in any robust sense. To say that the counter-facts determine whether a conditional is true or false is not to say that the counter-facts are context- or mind-independent features of reality. Just as one may be a moral subjectivist and hold that evaluative claims are true or false in virtue of subjective features of the agent making them, so too one can hold that conditionals are made true in part by, say, the agent's epistemic policies (indeed this one natural interpretation of the role of suppositions in determining the counter-facts). It is true that many non-factualists want to deny that conditionals have any truth-values at all,

[16] If we so choose, we can apply the formal model in such a way as to map facts onto what we consider to be real or observable or decidable, and counter-facts onto what is not. But this may not correspond very clearly with more usual ways of distinguishing between factual and modal properties.

agent dependent or otherwise. But I suspect that the main reason for this is the belief that the triviality results show that by reaping the advantages of a truth-compositional semantics they must deny themselves the explanatory resource of Adams's Thesis. But in this regard, as I will now show, they are wrong.

5. Probabilities of Conditionals

I have argued that in order to represent the different uncertainties associated with suppositions we need not one probability measure, but many: one for each supposition in fact. To examine the implications for conditionals, let *p* be a probability mass function on W that measures the probability that any world is the actual one and *q* be a probability mass function on A that measures the probability that any world is the counter-actual one, on the supposition that *A*. (I will say more about the interpretation of this probabilities later on). Finally let *pr* be a joint probability mass function on the pairs of worlds that lie in the table cells, measuring the joint probabilities of actuality and counter-actuality under the supposition that *A*. For example, $pr(<w_i,w_j>)$ is the probability that w_i is the actual world and w_j the counter-actual world on the supposition that *A*.

Let me say a word about joint probabilities. For *pr* to be the joint probability formed from *p* and *q* - i.e. if *p* and *q* are the marginal probabilities of the joint probability *pr* - it must be the case that *pr* is defined on the product domain W×A and that the three probability measures are related by the following condition:

Marginalisation:

$$\sum_{w_i \in W} pr(< w_i, w_j >) = q(w_j)$$
$$\sum_{w_j \in A} pr(< w_i, w_j >) = p(w_i)$$

Let $pr(\bullet|w_i)$ be the conditional probability mass function on W_A, given that w_i is the actual world. It then follows from the Marginalisation property that:

$$pr(< w_i, w_j >) = p(w_i) \times pr(w_j \mid w_i)$$

where $pr(w_j|w_i)$ is the probability that w_j is the counter-actual A-world given that w_i is the actual world. Now it follows from Centring that:

$$pr(w_1 \mid w_1) = 1$$
$$pr(w_2 \mid w_1) = 0$$

Hence our total state of uncertainty can be summarised as follows:

	Supposed A-Worlds	
Possible Worlds	**W1**	**W2**
W1	*p*(w1)	0
W2	0	*p*(w2)
W3	*p*(w3).*pr*(w1|w3)	*p*(w3).*pr*(w2|w3)
W4	*p*(w4).*pr*(w1|w4)	*p*(w4).*pr*(w2|w4)

This representation of our state of uncertainty in conjunction with the earlier claim that the semantic contents of conditionals are given by columns of the table ensures satisfaction of a probabilistic version of the Ramsey Test hypothesis. For instance, the conditional sentence $A \rightarrow B$ has as its truth conditions the W1 column of the table. This column has probability q(w1), which is the probability that w1 is the counter-actual A-world. Since w1 is the only counter-actual A-world at which *B* is true, it follows that the probability of $A{\rightarrow}B$ *must* equal the probability of B on the supposition that A. This is not just a feature of our very simple example, but is intrinsic to the way in which the multi-dimensional possible-worlds models being advocated here are constructed. Indeed, it is reasonable to say that such models encode the Ramsey Test.

To derive more specific versions of the hypothesis, appropriate for particular modes of supposition, further constraints need to be placed on the relation between marginal probabilities. Three candidates for characterising evidential supposition are salient:

1. *Stochastic Independence*: For all $w_i \in W$ and $w_j \in A$:

$$pr(w_j \mid w_i) = q(w_j)$$

2. *Counterfactual Independence*: For all $w_i \notin A$ and $w_j \in A$:

$$pr(w_j \mid w_i) = q(w_j)$$

3. *Restricted Independence*: For all $w_j \in A$:

$$pr(w_j \mid A) = q(w_j)$$

The three conditions are in decreasing order of strength. The most demanding, Stochastic Independence, says that the probability of a world being the counter-actual A-world is independent of what world is the actual world. Although it has the great virtue of considerably simplifying the assignment of probabilities, it is in conflict with Centring. Not so Counterfactual Independence, which says that the probability of a world being the counter-actual A-world is independent of any world at which A is false. This condition is of interest primarily because it is the counterpart in the multi-dimensional possible-worlds space of McGee's Independence condition on sentences. It is notable that Counterfactual Independence, together with Centring, is sufficient for the joint probabilities of world-pairs to be completely determined by the marginal measures *p* and *q*. For instance in our four-world example, the relevant uncertainties would now be as follows.

	Supposed A-Worlds	
Possible Worlds	**W1**	**W2**
W1	p(w1).1	0
W2	0	p(w2).1
W3	p(w3).q(w1)	p(w3).q(w2)
W4	p(w4).q(w1)	p(w4).q(w2)

Counterfactual Independence is subject to the same counter-examples as McGee's Independence condition (which it implies). Not so the third of the conditions, Restricted Independence, which says that the probability of a world being the counter-actual A-world is independent of the truth or falsity of *A* (this is the counterpart of our eponymous condition on sentences). Though it is the weakest of the three, it has two notable consequences: (i) it allows reduction of uncertainty regarding counter-actuality to ordinary conditional uncertainty about actuality, and (ii) it ensures that Adams's Thesis holds. To see this, note that in virtue of Centring, $pr(w_i,w_j) = p(w_i)$ and hence that it follows from Restricted Independence and Marginalisation that for all $w_j \in W_A$:

$$q(w_j) = pr(w_j \mid \mathrm{A}) = \frac{p(w_j)}{P(\mathrm{A})} = p(w_j \mid \mathrm{A})$$

So q is $p(\bullet|A)$. Now on our assignment of semantic values to conditionals the probability of $A \rightarrow B$ is given by $q(w_1)$, which we have just proven to equal $p(w_1|A)$. And this is just $P(B|A)$, in accordance with Adams's Thesis.

Under what conditions should one expect the Restricted Independence condition to hold? Answer: whenever the conditional under consideration is an indicative conditional. For, as we argued before, the mode of supposition that is relevant to indicative conditionals is evidential and standard theories of evidential supposition require that the probability of an event under the supposition that some condition is, as a matter of fact, true is independent of the probability of the condition itself (in Bayesian theory this is known as the Rigidity assumption). Given this, it is an immediate consequence of the treatment of semantic content and uncertainty presented here that the probability of truth of a simple indicative conditional is indeed just the conditional probability of its consequent, given the truth of its antecedent. So Adams's Thesis is not only consistent with a truth-conditional semantics proposed here but, given a plausible view about evidential supposition, is required by it.

6. Expected Truth and Chance

With the derivation of Adams's Thesis the main task of the paper is completed. But before concluding let me make some remarks on the connection between the proposed theory and some of the others we have looked at. Firstly, our model of uncertainty can be used to derive the account of conditionals in which the contents of sentences are treated as random variables whose values are the expected truth values of the sentences. To do so we simply need to take the semantic value of a sentence at a world w to be given by the sum of conditional probabilities, given w, of the counter-actual worlds at which the sentence is true. Thus, in our simple model, the random variable associated with conditional $A \rightarrow B$ would take the value $pr(w1|w_i)$ at each world w_i; i.e. 1 at w1, 0 at w2, $pr(w1|w3)$ at w3 and $pr(w1|w4)$ at w4. (Evidently more assumptions are required in order that these values agree with the JS-semantics). In essence then the random variable account can be thought of as a *projection*, onto a single dimension, of a two dimensional possible-worlds semantics, in which the probabilities of truth of pairs of worlds imply probabilities of expected truth at singleton worlds.

Secondly, construing the semantic values of sentences at worlds as expected truth-values is consistent with a variety of views regarding the nature of these expectations and the extent to which they are constrained by the facts. On McGee's account they depend, at least partially, on pragmatic considerations encoded in personal probabilities. But more objective interpretations of the uncertainty associated with counter-facts may be appropriate. For instance, when the chances (or relative frequencies or propensities) of the relevant events are given at a possible world w_i, then the mass function $pr(\bullet|w_i)$ could be construed as a measure of the conditional chances (or frequencies or propensities) given A, at w_i. More precisely, it could be construed as a measure of the degrees of belief in the counter-facts rationally required by the given objective probabilities. Doing so would allow us to say that even if the counter-facts are not completely determined by the facts, our expectations regarding them are. And hence that the facts suffice to determine the expected truth values of conditionals.

For definiteness, let's focus on the interpretation of the objective uncertainties as conditional chances and call it the Chance view. On this view, expected truth values are fixed by conditional chances and our subjective uncertainty concerns only what the facts about chances are. For instance, in any world in which the coin is fair, the fact that it has an even chance of landing heads or tails implies that the expected truth value of the sentence $T\rightarrow H$ (that the coin will or would land heads if it is or were tossed) must equal one-half. Similarly in worlds in which it is biased towards heads (or tails) the expected truth value of the $T\rightarrow H$ must be one (or zero). In general then our uncertainty about the truth value of $T\rightarrow H$ decomposes into two components: subjective uncertainty about what, as a matter of fact, the chances are, and objective uncertainty regarding the truth of $T\rightarrow H$, given the chances. (Note that since chances are time-dependent, it will follow that expected truth values are too.)

The Chance view has some notable consequences. Firstly, it leads to the 'right' kind of violation of McGee's principle of Counterfactual Independence. Recall that this principle implies that the expected truth value of a conditional sentence is the same at every world at which its antecedent is false. But although this was true in the first coin-tossing example we considered, it was not the second: for instance the expected truth value of $T\rightarrow H$ in this case was 1 at w2, but 0 at w4. This can now be explained by the fact that in the second example, unlike the first, the conditional chances are not invariant across worlds. Secondly, the Chance view implies the version of the Ramsey Test hypothesis that I previously labelled Skyrms's Thesis, namely that the probability of a conditional $A\rightarrow B$ equals the expected chance of B given A. This follows, by application of equation (2), from the identification of the semantic value of $A \rightarrow B$ at any world w with its expected truth value at w, and this in turn with the conditional chances at w of B, given A.

Despite these attractive features, the Chance view has some problems. Firstly, it conflicts with Centring, which requires that the expected truth value of $A \rightarrow B$ be 1 at any world in which both A and B are true, whatever the conditional chances of B given A at the world. Secondly, Skyrms's Thesis was intended as an account of counterfactuals only, so any principle that implies that it holds for all conditionals must be too strong. And, indeed, a weaker constraint on the relation between chances and expected truth values suggests itself. Let $\{Ch_i\}$ be a partition of the space of possible worlds such that the conditional chances of B given A at each of the worlds in any cell are the same. Then the *average* of the expected truth values of a simple conditional $A \rightarrow B$ at the worlds within a cell should equal the conditional chances of B given A at these worlds. Formally, let Q be a probability measure on the counter-factual A-events (subsets of W_A) and let $Ch_i(B|A)$ be the conditional chances of B given A according to hypothesis Ch_i. Then, by what I am tempted to call the Principal Suppositional Principle:

$$(PSP)\ Q(B \mid Ch_i) = Ch_i(B \mid A)$$

PSP say that the probability that if A were the case then B would be, given chance hypothesis Ch_i, should equal the chance of B given that A on this hypothesis. This weaker constraint is consistent with Centring and still conflicts with Counterfactual Independence in the right kind of way. Furthermore it is consistent with both Adams's and Skyrms's Thesis. To see this note that by the Ramsey Test hypothesis, of which they are both instances, the probability of $A \rightarrow B$ equals $Q(B)$. And by PSP:

$$Q(B) = \sum_i Ch_i(B \mid A).Q(Ch_i)$$

Then to get our two theses we must require that $Q(Ch_i)$ equal to $P(Ch_i \mid A)$ if supposition is evidential and equal to $P(Ch_i)$ if it is counterfactual. All of this makes PSP very attractive as a general hypothesis about the relation between chances and expected truth values. But to work out the details carefully would require considerable embellishment of our semantic model, and so for now the topic will have to be placed under the label of future research.

7. Concluding Remarks

The crucial modification to standard theory proposed in this paper is the representation of semantic content by ordered sets of possible worlds rather than just sets of worlds. Up to this point we have dealt only with a very simple example involving just one supposition, but the treatment can be generalised in a straightforward way. For simplicity we will assume a given set W of possible worlds and an associated assignment of truth-values to factual sentences and then extend the assignment to conditionals using the notion of a suppositional or multidimensional possible-worlds space.

Let L be a language containing only factual sentences, but closed under conjunction, disjunction and negation. And as before let v be an orthodox assignment of truth values to L-sentences, with $v_w(A)$ denoting the truth value of sentence A at world w and $[A]$ the set of worlds at which A is true. Now let SL be a simple conditional language closed under conjunction, disjunction and negation and such that sentences A and B belong to L iff sentence $A \rightarrow B$ belongs to SL.

A suppositional space S is defined formally as follows. Let $\wp(W) = \{X_i\}$ be the set of all subsets of W. Each member of this set, X_i, is a possible factual supposition. S is now the space spanned by all the possible suppositions i.e. $S = W \times X_1 \times \ldots \times X_n$. Any element of S is a vector $\omega = \langle w, w_{X1}, \ldots, w_{Xn} \rangle$ of worlds with w_0 being a possible actual world and each w_{Xi} being a possible counter-actual world under the supposition that X_i. For any vector ω and subset X_i, of W, let $\omega^{(X_i)}$ denote w_{Xi}, the component of ω corresponding to the supposition that X_i. With this modification to the possible-worlds framework in place, we can state corresponding versions of the four propositions characterising the orthodoxy.

Semantics: An interpretation of language SL is a mapping v^* from pairs of SL-sentences and ordered sets of possible worlds to semantic values satisfying the conditions of (1a*) Bivalence and (1b*) Boolean Composition[17], and such that for any factual sentence B and world vector $\omega = \langle w_0, \ldots, w_n \rangle$, $v^*_\omega(B) = v_w(B)$ and:

(1c*) Conditionals: $v^*_\omega(A \rightarrow B) = v_{\omega([A])}(B)$

Pragmatics: Let pr be a joint probability mass function on the set of world vectors, with $pr(\omega)$ measuring the probability that w_0 is the actual world and w_i the counter-actual X_i world. Then rational degrees of belief in L-sentences are measured by a probability function Pr such that for all L-sentences A:

(2*) $$\Pr(A) = \mathrm{E}(v^*(A)) = \sum_{\omega \in S} v^*_\omega(A).pr(\omega)$$

[17] A starred condition is the same as the original orthodox condition with any reference to a world w being replaced by a reference to a vector ω.

Logic: Let [A]* denote the content of the sentence A, i.e. the sets of world vectors making it true. Then:

(3*) $A \rangle B$ iff $[A]^* \subseteq [B]^*$

Explanation: Let $\Pi = \{pr_j\}$ be the set of all permissible joint probability mass functions on the set of world vectors S and $V_L = \{v^*_i\}$ be the set of all permissible assignments of semantic values to sentence of L. A multidimensional possible-worlds model (MPW-model for short) of L is a structure $<v^*, pr>$ where v^* belongs to V_L and p to Π. Then

(4*) Every pair $< v^*, pr >$ in $V_L \times \Pi$ is an MPW-model.

The semantic model is easily generalised to handle languages containing nested conditionals. To do so we need to define a hierarchy of suppositional spaces, each level consisting of vectors of elements of the preceding (lower) level, with each component of any vector corresponding to a supposition. Then we extend the semantic assignment from simple conditionals to more complicated ones in just the same way as we extended the semantic assignment from factual sentences to simple conditionals. In fact the procedure is identical from a formal point of view since, from the perspective of each level in the hierarchy, the elements of the preceding (lower) one can be construed as worlds and any sets of these worlds as a suppositions. Once this has been done, we can ask questions about how the hierarchical structure of the suppositional space might constrain the nesting of conditionals both semantically and pragmatically.[18] But since no claims about nested

[18] For instance what properties would suffice to require conditionals to respect the import-export condition.

conditionals have been explored in this paper, I will not spell this generalisation out any further.

Let us return to the problem that we began with. The permissibility of an assignment of semantic values to sentences is determined by the semantic conditions 1a*, 1b* and 1c*. On the other hand, which joint probabilities are permissible depends on the constraints imposed on the relation between joint and marginal probabilities. And these in turn depend on the kind of supposition being modelled; in particular whether it is of the evidential or counterfactual variety. In the case of evidential supposition, I have argued that the condition of Restricted Independence is appropriate, a condition which, jointly with (1c*) and (2*), implies Adams's Thesis.

In the light of the similarity between the multidimensional possible-worlds models and the orthodox uni-dimensional ones, one may wonder if the triviality results for Adams's Thesis do not still apply in the new modified framework. Constructing a triviality argument seems simple enough. Let any occurrence of the word 'possible world' in your favourite triviality result be replaced by the phrase 'vector of possible worlds' and you appear to get a triviality result for the theory presented here.

In fact this is not so, for one crucial assumption is no longer satisfied. Condition (4*), unlike the original condition (4), allows for restrictions on permissible belief measures above and beyond the requirement that they be probabilities. In particular, it allows for restrictions on the relation between the joint probabilities on the suppositional space (which determine beliefs in sentences) and marginal probabilities defined on spaces of possible worlds, such as those contained in the various independence principles canvassed above. As such they cannot be framed without the additional structure contained in suppositional spaces, and so cannot be stated within an orthodox possible world model. Crucially, amongst those belief measures ruled out by the independence constraints are just those required for the triviality results to go through. For example, given the Restricted Independence condition, there can be no joint probability *pr* on W×A with marginals *p* on W and $q = p(.|A)$ on A that does not satisfy the Preservation condition. This is because if $P(B) = 0$ then $Q(B) = P(B|A) = 0$.[19] So the non-trivial accommodation of the Preservation condition, and indeed of Adams's Thesis, within a modified possible-worlds framework is assured.

References

Adams, Ernest. 1975. *The Logic of Conditionals*. Dordrecht and Boston: Reidel.

Bradley, Richard. 2000. "A Preservation Condition for Conditionals." *Analysis* 60: 219-22.

————. 2002. "Indicative Conditionals." *Erkenntnis* 56: 345-378

————. 2011. "Conditionals and Supposition-based Reasoning." *Topoi* 30: 39-45

Döring, Frank. 1994. "Probabilities of Conditionals." *The Philosophical Review* 103: 689-700.

Douven, Igor. 2007. "On Bradley's Preservation Condition for Conditionals." *Erkenntnis* 67: 111-118.

Edgington, Dorothy. 1991. "Do Conditionals Have Truth-Conditions?" In *Conditionals*, ed. Frank Jackson, 176-201. Oxford: Oxford University Press

[19] It also has the implication that conditionalisation is not the correct way to revise degrees of belief in conditionals, thereby blocking Lewis's (1965) second triviality result. For a full

————. 1995. “On Conditionals.” *Mind* 104: 235-329.

Eels, Ellery and Brian Skyrms. eds. 1995. *Probabilities and Conditionals: Belief Revision and Rational Decision*. Cambridge: Cambridge University Press.

Gibbard, Alan. 1981: "Two Recent Theories of Conditionals." In *Ifs: Conditionals, Belief, Decision, Chance and Time*, ed. William L. Harper, Robert Stalnaker and Glenn Pearce, 211-247. Dordrecht: Reidel.

Hájek, Alan. 1989. “Probabilities of Conditionals - Revisited.” *Journal of Philosophical Logic* 18: 423-428.

Jackson, Frank. 1979. "On Assertion and Indicative Conditionals." *The Philosophical Review* 88: 565-89.

Jeffrey, Richard C. 1964. “If (abstract).” *Journal of Philosophy* 61: 702-703.

————. 1991. “Matter-of-fact Conditionals” *Aristotelian Society Supplementary Volume* 65: 161-183.

Jeffrey, Richard C., and Robert Stalnaker. 1994. “Conditionals as Random Variables.” In *Probabilities and Conditionals*, ed. Ellery Eels and Brian Skyrms, 31-46. Cambridge: Cambridge University Press.

Kaufmann, Stefan. 2005 “Conditional predictions: A probabilistic account.” *Linguistics and Philosophy* 28:181-231.

Lance, Mark. 1991. "Probabilistic Dependence among Conditionals." *Philosophical Review* 100: 269-76.

Lewis, David K. 1973. *Counterfactuals*. Oxford and Malden, MA: Basil Blackwell

————. 1976. “Probabilities of Conditionals and Conditional Probabilities.” *Philosophical Review* 85: 297-315.

discussion see Bradley 2011.

McGee, Vann. (1989). “Conditional Probabilities and Compounds of Conditionals.” *Philosophical Review* 48: 485-541.

McDermott, Michael. 1996. “On the Truth Conditions of Certain ‘If’-Sentences”. *Philosophical Review* 105: 1-37.

Milne, Peter. 1997. "Bruno de Finetti and the Logic of Conditional Events." *British Journal for the Philosophy of Science* 48: 195-232.

Over, David, and Jonathan Evans. 2003. “The Probability of Conditionals: The Psychological Evidence.” *Mind and Language* 18: 340–358.

Skyrms, Brian. 1981. “The Prior Propensity Account of Subjunctive Conditionals.” In *Ifs: Conditionals, Belief, Decision, Chance and Time*, ed. William L. Harper, Robert Stalnaker and Glenn Pearce, 259-265. Dordrecht: Reidel.

————. 1994. “Adams Conditionals.” In *Probabilities and Conditionals*, ed. Ellery Eels and Brian Skyrms, 31-46. Cambridge: Cambridge University Press.

Stalnaker, Robert. 1968. “A Theory of Conditionals.” In *Studies in Logical Theory (American Philosophical Quarterly Monograph)*, ed. Nicholas Rescher, 98-112. Oxford: Basil Blackwell.

————. 1970. “Probability and Conditionals.” *Philosophy of Science* 37: 64-80.

Van Fraassen, Bas. 1975. “Probabilities of Conditionals.” In *Foundations of Probability and Statistics*, ed. William Harper and C.A. Hooker, vol. l, 261-308. Reidel: Dordrecht