



OXFORD JOURNALS  
OXFORD UNIVERSITY PRESS

ANALYSIS

---

Classifying Conditionals

Author(s): Frank Jackson

Source: *Analysis*, Vol. 50, No. 2 (Mar., 1990), pp. 134-147

Published by: Oxford University Press on behalf of The Analysis Committee

Stable URL: <http://www.jstor.org/stable/3328856>

Accessed: 14/06/2010 12:34

---

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=oup>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).



Oxford University Press and The Analysis Committee are collaborating with JSTOR to digitize, preserve and extend access to *Analysis*.

<http://www.jstor.org>

## CLASSIFYING CONDITIONALS

By FRANK JACKSON

### CONSIDER

- (1) If Booth had not killed Lincoln, someone else would have
- (2) If Booth does not kill Lincoln, someone else will  
and
- (3) If Booth did not kill Lincoln, someone else did.

Many writers agree that (1) is importantly different from (3). It is not just that we deny (1) and affirm (3) – a matter everyone agrees about – rather the explanation of why we deny (1) and affirm (3) is to be found in an important semantic difference between (1) and (3).<sup>1</sup> The details vary (needless to say), but there is considerable agreement that the key to the meaning of (1) lies in the fact that (1)'s truth conditions are to be given in terms of possible worlds in the style of Robert Stalnaker [13], or of David Lewis [11], or of some reasonable variant thereon; whereas the key to the meaning of (3) lies in the fact that the justified assertibility of (3) is given by the conditional probability of (3)'s consequent given its antecedent, or in terms of some reasonable variant thereon. (See e.g. Ernest Adams [1], [2].) This approach to (3)'s meaning is sometimes associated with the doctrine that (3) does not have truth conditions, and sometimes with the doctrine that (3) has the truth conditions of the material conditional.<sup>2</sup>

This paper takes for granted this general attitude to the relationship between (1) and (3). Our question is what then to say about (2). Should we group (2) with (1), or with (3)? If – as is common – we describe (1) as a subjunctive conditional and (3) as an indicative conditional, and go on to express the generalization of the idea that (1) and (3) are semantically distinct in an important way by saying that subjunctive and indicative conditionals are importantly semantically distinct, we presumably place (2) with (3). For (2) would commonly be classified as indicative. But this would be weak *qua* argument for classifying (2) with (3). (As has been widely noted, recently by Jonathan Bennett [3] and Michael Pendelbury [12].) The fact that (1) is subjunctive whereas (3) is indicative is far from the only difference between them which might reasonably be supposed to be the relevant one, particularly in view of the notorious elusiveness of the subjunctive/

<sup>1</sup> See e.g. Jonathan Bennett [3], Allan Gibbard [8], David Lewis [11] and Frank Jackson [10]. For a dissenting voice see Brian Ellis [6]. I follow Bennett in using Booth and Lincoln instead of Oswald and Kennedy. You are of course supposed to make a 'Warrenite' assumption about the example.

<sup>2</sup> For examples of the former response see Adams [2], and Dorothy Edgington [5]; for an example of the latter response see Jackson [10].

indicative distinction in English. Moreover, one of the main reasons for holding that (1) and (3) are semantically distinct is the noted point that we deny (1) and affirm (3); but if we go by denial and affirmation, then (2) most naturally goes with (1). As Brian Ellis in [6] and [7] has particularly emphasized, our reasons for denying (1) – the absence of a backup killer, for instance – are those which would justify denying (2) from a temporal perspective just prior to Booth shooting. What is more, suppose we acquired reason to affirm (1) by learning that Booth's wife was ready to fire from the stalls if he missed from the balcony; we would in that case hold that (2) was the right thing to say beforehand, and so going by affirmation and denial would in this case as well put (1) and (2) together.

However, the evidence from affirmation and denial for classing (1) and (2) together is not as strong as it might seem at first. There is an important difference between the data about affirmation and denial as it applies to separating (1) from (3), and as it applies to where to locate (2). From a given epistemological standpoint – roughly, that which we actually occupy – we deny (1) and affirm (3). That is the datum that most often motivates giving different accounts of (1) and (3). But it is not true that from a given epistemological standpoint we deny (2) while denying (1) and affirming (3). We deny (1) and affirm (3) from the standpoint we actually occupy after the event; but deny (2) from a standpoint we imagine ourselves to be occupying beforehand. There is no one epistemological standpoint from which we give different answers about (2) and (3). Hence, in the case of (2) and (3) it is an open possibility that the difference in regard to affirmation and denial can be explained by the difference in epistemological standpoints, thus allowing us to class (2) and (3) together, semantically speaking, and apart from (1).

The principal burden of this paper is that we should, semantically speaking, class (2) and (3) together, and apart from (1). The semantic line goes where the mood line goes, between (1) and (2), not between (2) and (3). After I have argued this, I will return to the question of providing an epistemological explanation of why we affirm (3) but from the imagined perspective before the event, deny (2).

I think that there are a number of reasons for classifying (2) with (3) (I give some of these reasons in [10], pp. 40f), but the case I will develop here turns on that provided by Allan Gibbard's Sly Pete example in [8]. I think that the significance of this much discussed example has been left rather obscure in the literature. Indeed, Bennett takes the example to provide reason for classifying (2) with (1).<sup>3</sup> I will argue, though, that once the lesson of the

<sup>3</sup> See Bennett [3]. Michael Pendelbury [12] also holds (but for different reasons from Bennett's) that (2) should be classified with (1).

example is properly understood it provides a strong reason for classifying (2) with (3).

It will be useful to have labels for conditionals relevantly like (1), like (2), and like (3). I will call conditionals relevantly like (1) 'past subjunctives'; like (2) 'future indicatives'; and like (3) 'past indicatives'. I use these labels because they are convenient — you will know the conditionals I have in mind — and set to one side whether there are better labels to be had. Thus the main issue we will be concerned with can be expressed as how to locate semantically future indicative conditionals: with past subjunctives as Bennett holds, or with past indicative conditionals as I hold.<sup>4</sup>

## I THE NATURAL APPROACH TO CONDITIONALS

My argument will be that although the natural approach to conditionals applies to past subjunctive conditionals, the Sly Pete example shows that it does not apply to past indicative conditionals and does not apply to future indicative conditionals. And, as will be clear when I describe the natural approach, this difference between past subjunctive on the one hand, and past and future indicative conditionals on the other, is most certainly a semantic one.

The natural approach is, as its name suggests, the natural one to take, the one we would like to take if we could. It has two clauses: (Support) and (Conditional non-contradiction).

(Support). On the face of it, acquiring information which justifies the assertion of a conditional is not different in kind from acquiring evidence which justifies the assertion of a non-conditional sentence. Reading the weather report justifies my saying that it will rain, reading the rules of cricket justifies my saying that if it rained heavily, the match was cancelled; learning about the stock exchange may justify my saying that a certain stock will go up, or it may justify my saying that if a certain stock goes up, certain other stocks will go down, and so on and so forth. By saying that there is, on the face of it, not a difference in kind I simply mean that both cases seem alike in essentially involving the acquisition of evidence that makes what is asserted more likely to be true. The information is support in the standard sense of making very probably true. (Support) is thus incompatible with the view that conditionals do not have truth conditions, only assertion conditions. But I think that it is important to appreciate that this deflationary position is one we need to be driven into; it is not one to embrace early on. For it makes a mystery of the sense in which

<sup>4</sup> I am, of course, also disagreeing with Pendelbury [12]. But note that he labels (2) as well as (1) as 'subjunctive', see [12], p. 192. Hence, we come out in verbal agreement but remain in disagreement over the substantive question of (2)'s placement.

information justifies, in the relevant sense, asserting a conditional. For instance, a lot of evidence that in the past, economic problems have been followed by election reverses for the government justifies saying that if the economy continues to go badly, the government will be in trouble at the next election. But this is not justification in the sense of making the saying morally justified, or pragmatically advisable, or in accord with good etiquette. What then is the relevant sense of justification? The obvious answer is that it is justification in the sense of making probably true.

(Conditional non-contradiction). There is something very wrong with a given person saying both 'if *A* then *B*' and 'if *A* then not *B*', or at least there is when *A* is consistent. In tandem with this is the fact that if you say 'if *A* then *B*' and I say 'if *A* then not *B*', we count as disagreeing, or at least we do when *A* is consistent. (Reductio proofs in Logic and Mathematics show that when *A* is inconsistent the situation is different.) The obvious explanation of these facts is that 'if *A* then *B*' and 'if *A* then not *B*' are (logical) contraries when *A* is consistent. We will follow the practice of referring to this as (Conditional non-contradiction). (The name perhaps suggests the doctrine that, when *A* is consistent, not only are 'if *A* then *B*' and 'if *A* then not *B*' contraries, they are also contradictories in the sense that they must have opposite truth values. The argument that follows goes through on either interpretation, but we will work in terms of the weaker one.)

If a possible worlds semantics of one of the more familiar kinds applies to a given class of conditionals, then there is no problem about holding that the natural approach applies to that class. If 'if *A* then *B*' is true if and only if (something along the lines of) the closest *A*-worlds are all *B*-worlds, then 'if *A* then *B*' is (a) truth-valued, (b) supported by evidence in the standard sense because information garnered about how things actually are can make it more or less likely that nearby antecedent worlds are consequent worlds, and (c) inconsistent with 'if *A* then not *B*' when *A* is consistent, for the closest *A*-worlds cannot be both *B*-worlds and not-*B*-worlds. I think, in common with many, that the possible worlds semantics applies to past subjunctive conditionals. Hence, I take it that there is no problem about adopting the natural approach for past subjunctive conditionals, including of course (1).

The problem is with past and future indicative conditionals. The Sly Pete example shows that (unfortunately) the natural approach fails for them; although, as we will see, it takes a little work to see exactly how to argue the point.

## II THE SLY PETE EXAMPLE AND SOME FALSE STARTS

There are a number of ways of setting out this example. We will work with the following version. It is in terms of past indicatives. We will discuss future indicatives shortly.

Sly Pete is playing poker. He has to decide between calling and folding. Informant *A* knows that Sly Pete is cheating, and that he is doing so by knowing what is in his opponent's hand. (Perhaps Sly Pete can see his opponent's hand in a mirror in the ceiling.) Informant *B* has surreptitiously observed both hands and knows that Sly Pete has the weaker hand. We can suppose that it is the last hand for the night in order to avoid the complication that Pete might want to lose the hand in order to encourage his opponent to bet higher later in the game. Informants *A* and *B* leave the room just before Pete makes his decision and report to you as follows. *A* says that if Pete called, he won. *B* says that if Pete called, he lost. You respond (correctly) by inferring that Pete did not call, he folded.

How does this example make trouble for the natural approach as applied to past indicative conditionals? Gibbard sees things this way. He points out that both informant *A* and informant *B* are fully justified in their assertions, are working from information properly so called, that is, which contains no mistakes, and are being sincere. From this he argues that 'we can see that neither is asserting anything false', and so that *if* 'both ... utterances express propositions, ... both express true propositions'.<sup>5</sup> To put matters in our terms, Gibbard is arguing that the (Support) part of the natural approach when combined with the fact that informant *A* and informant *B* are sincere and are proceeding in a fully justified manner from data that contains no mistakes would mean that what they each say is true, and so that (Conditional non-contradiction) fails. Hence, the natural approach is false, for its two clauses cannot be true together.

The trouble with this argument is that it is perfectly possible to proceed in a fully justified manner from information properly so called to a false conclusion. That is the point behind the familiar distinction between deductive and inductive arguments. In particular, it is perfectly possible for two different bodies of information relevant to a common subject matter to sustain incompatible conclusions about that subject matter. Smith knows that Petersen is a churchgoing Swede and that nearly all churchgoing Swedes are Protestants. She is entitled to infer with considerable confidence that Petersen is a Protestant. Jones knows that Petersen has Catholic parents and that nearly all churchgoers with Catholic parents are Catholic. She is entitled to infer with considerable confidence that Petersen is a Catholic. All this is compatible with Smith's information being information properly speaking, correct

<sup>5</sup> Gibbard [8], p. 231. Helpful correspondence with Gibbard has convinced me that though my interpretation of Gibbard's argument is the best interpretation of the printed word, it is probably not the line of argument he intended, although I confess that the precise direction of the intended line of argument is still obscure to me.

in every detail, and likewise for Jones's information. Of course, their two bodies of information are incomplete, but the same goes for the bodies of information available to informant *A* and to informant *B* in the Sly Pete example.

There is, of course, a special feature of the Sly Pete example. In addition to the fact that both informants proceed justifiably from correct if incomplete information, there is the point that you, the hearer, are able to use their two utterances constructively to infer that Sly Pete did not call. This suggests that (a) 'if Sly Pete called, he won' conjoined with 'if Sly Pete called, he lost' entails that Sly Pete did not call, and that (b) it is rational to accept, in the circumstances as specified in the Sly Pete example, the conjunction of 'if Sly Pete called, he won' with 'if Sly Pete called, he lost'.

This way of construing the argument based on the Sly Pete example avoids the mistake of supposing that what comes rationally from what is the case must be true. It rests on the idea that if it is rational to accept 'if Sly Pete called, he won' together with 'if Sly Pete called, he lost', they can hardly be contraries.<sup>6</sup> The problem for this way of construing the argument is that it is not sufficiently obvious that our inference to 'Sly Pete did not call' depends *inter alia* on accepting the conjunction of the two conditionals. There is another hypothesis available to explain what is going on.

An important feature of the Sly Pete story is that you know the kind of information available to your two informants. You know that one informant knows that Sly Pete is cheating by knowing his opponent's hand, and that the other informant knows that Sly Pete's opponent has the stronger hand; and this fact plays a crucial role in warranting your inference from their two utterances to the conclusion that Sly Pete did not call. For suppose that you had had no idea of the evidence available to them, and that all you knew was that they were sincere and sensible. In that case you could have done nothing constructive with their two utterances, and in particular you could not have inferred that Sly Pete did not call. For suppose you cannot now remember whether you had your child inoculated against whooping cough, and one doctor says to you 'if you had your child inoculated against whooping cough, then you did the right thing', while another says to you 'if you had your child inoculated, you did the wrong thing'; you do not infer without further ado that you did not have your child inoculated against whooping cough. If in a sporting broadcast you hear Fred Stolle say that if Chris Evert won last night, she went one ahead in the head to head tally against Martina Navratilova,

<sup>6</sup> If *S* and *S\** are sentences which are hard to comprehend fully, it can be rational to accept *S* along with *S\** even when *S* and *S\** are (unobviously) contraries. This is a point familiar to logicians and mathematicians. But it is not plausible to think that our two sentences fall into this category.

and then hear John Newcombe say that if Evert won, the tally is dead level, you do not take your high regard for the sincerity and good sense of Stolle and Newcombe as justifying the conclusion that Chris Evert lost. You simply infer that one of them has made a mistake.

Now, if it is crucial in the Sly Pete example that you know a good deal about your informant's evidence, there is an alternative hypothesis about how you reached your conclusion that Sly Pete did not call. We do not have to suppose that you inferred from the two conditional utterances taken together — a supposition which would indeed count against (Conditional non-contradiction) — we can instead suppose that you inferred from what you know of the evidence available to your informants. You know that informant *A* is using evidence that Sly Pete is cheating, informant *B* is using evidence that Sly Pete's hand is the weaker, and those two bits of evidence taken together are in themselves enough in the circumstances to warrant inferring that he did not call. We do not need to go via the informants' conditional utterances to explain how you were entitled to your conclusion.

In a variant on the Sly Pete example your informants' conditional utterances tell you something important about their respective bodies of evidence. In the variant you do not know initially that informant *A* knows that Sly Pete is cheating. You know merely that informant *A* knows whether or not Sly Pete is cheating, and knows nothing else relevant. Similarly, you do not know initially that informant *B* knows that Sly Pete has the weaker hand. You know merely that informant *B* knows who has the weaker hand, and knows nothing else relevant. In this variant case the conditional utterances of *A* and *B* do give you highly relevant information. When *A* says 'if Sly Pete called, he won', you learn that your informant knows that Sly Pete is cheating, for otherwise, given what you already know about *A*, *A* would not be justified in saying that if Sly Pete called, he won. When *B* says 'if Sly Pete called, he lost', you learn that your informant knows that Sly Pete's hand is the weaker one, for otherwise, given what you already know about *B*, *B* would not be justified in saying that if Sly Pete called, he lost. (We assume known sincerity and good sense in *A* and in *B*.) But even in this version of the Sly Pete example where your two informants' utterances play a major role in enabling you to conclude that Sly Pete did not call, we do not have to suppose that it is what they say as such that you are using. A believer in (Conditional non-contradiction) can insist that you are using what their two utterances tell you about their evidence, not the contents of the utterances themselves, to conclude that Sly Pete did not call. And this insistence would not be *ad hoc*. Believers in (Conditional non-contradiction) are entitled to make much of the point already emphasized that the two assertions of your informants, even if known to be sincere and sensible, in the absence of any informa-



tion about the evidence that lies behind them, would not enable you to conclude that Sly Pete did not call. How so, if the evidence that backs their assertions is not the key?

### III THE SLY PETE ARGUMENT AND A PRINCIPLE ABOUT EVIDENCE

Despite these failures, I think that it is possible to mount an argument from the Sly Pete example against the natural approach. I will refer to such an argument as a Sly Pete argument. It rests on a principle about evidence in favour of hypotheses which are logical contraries.

You might be tempted by the following principle about evidence in favour of contraries: if  $e_1$  strongly supports  $H_1$ ,  $e_2$  strongly supports  $H_2$ , and  $H_1$  and  $H_2$  are contraries, then  $e_1$  is unlikely given  $e_2$ , and  $e_2$  is unlikely given  $e_1$ . But there is a counterexample to it. Set  $e_1$  to be 'Petersen is a Swedish Catholic',  $e_2$  to be 'Petersen is a Swede',  $H_1$  to be 'Petersen is Catholic', and  $H_2$  to be 'Petersen is not Catholic'.  $e_1$  strongly supports  $H_1$  (by entailing it),  $e_2$  strongly supports  $H_2$  (relative to a background which includes the fact that nearly all Swedes are Protestants),  $H_1$  and  $H_2$  are contraries, and yet it is false that  $e_2$  is improbable given  $e_1$ . In fact  $e_1$  entails  $e_2$ . It is, however, true that  $e_1$  is improbable given  $e_2$  (relative to the background evidence), and indeed the following principle is valid.

- (P) If  $e_1$  is strong evidence for  $H_1$ ,  $e_2$  is strong evidence for  $H_2$ , and  $H_1$  and  $H_2$  are contraries, then  $e_1$  is improbable given  $e_2$  or  $e_2$  is improbable given  $e_1$  (or both).

More precisely, we can specify a value such that if  $\Pr(H_1/e_1)$  and  $\Pr(H_2/e_2)$  both exceed it, one or both of  $\Pr(e_1/e_2)$  and  $\Pr(e_2/e_1)$  are  $< 0.5$ .<sup>7</sup>

It is by reference to (P) that we show that the Sly Pete example refutes the natural approach to conditionals as applied to past indicatives. Informant *A*'s information is that Sly Pete knows what is in his opponent's hand. This information along with such facts as that Sly Pete wants not to lose money on this hand and understands the rules of poker, very strongly warrants *A*'s assertion that if Sly Pete called, he won. Similarly, informant *B*'s information that Sly Pete has the weaker hand, in the appropriate circumstances very strongly warrants *B*'s assertion that if Sly Pete called, he lost. However, Sly Pete's knowing his opponent's hand need not count against Sly Pete's having the weaker hand. Similarly, Sly Pete's having the weaker hand need not count against his knowing his

<sup>7</sup> One way to get the result is to start with the special case where  $H_1$  and  $H_2$  are contradictories so that, for instance,  $\Pr(e_1/e_2) < 0.5$  iff  $\Pr(e_1e_2H_1) + \Pr(e_1e_2H_2) < \Pr(e_2H_1 \sim e_1) + \Pr(e_2H_2 \sim e_1)$ .

opponent's hand. Typically, 'Sly Pete knows his opponent's hand' and 'Sly Pete has the weaker hand' are evidentially neutral with respect to each other. In the example as originally described by Gibbard, Sly Pete has an accomplice who signals the contents of the opponent's hand to Pete. We imagined that Pete can see his opponent's hand reflected in a mirror in the ceiling. Typically, the successful operation of systems of these kinds are probabilistically quite independent of which hand is the stronger of the two.

The same point can be made in terms of other examples. Information that a certain motion put at a certain meeting was a bad one would strongly warrant asserting 'if Fred voted for the motion, he did the wrong thing'. Information that Fred has excellent judgment in such matters would strongly warrant asserting 'if Fred voted for the motion, he did the right thing'. However, the fact that a certain motion was bad does in itself not count against (or for) the quality of Fred's judgment; and, conversely, Fred's excellent judgment does not in itself count against (or for) the motion being bad. We must conclude, therefore, that the Sly Pete example shows via (P) that the natural approach fails for past indicative conditionals. It cannot simultaneously be the case that (a) the two bodies of information that very strongly warrant the assertion, respectively, of 'if Sly Pete called, he won' and 'if Sly Pete called, he lost', are evidence in the standard sense of making the conditional whose assertion is warranted highly probably true, and (b) the conditionals are logical contraries. (a) and (b) cannot be true together because the two bodies of information may each be not improbable given the other. It follows that (Support) and (Conditional non-contradiction) clash in the case of past indicatives. But as we saw (or rather noted had been widely granted), the natural theory is true for past subjunctives, so we have the result that past indicative conditionals are semantically distinct from past subjunctive conditionals.

By itself this result is no advance. We *started* by noting the semantic distinctness of (1) from (3). We are now though in a position to argue that (2), the future indicative, should be classed with (3), the past indicative, rather than with (1), the past subjunctive. We can, now say with some precision how a Sly Pete argument can show that the natural approach to conditionals fails for conditionals of some specified class. It does so by describing a case where there are two bodies of evidence  $e_1$  and  $e_2$ , and a pair of conditionals 'if  $A$  then  $B$ ' and 'if  $A$  then not  $B$ ' of the class in question with  $A$  consistent, such that against some given background: (a)  $e_1$  by itself strongly warrants (as highly as you care to make it by varying the background) asserting 'if  $A$  then  $B$ ', (b)  $e_2$  by itself strongly warrants asserting 'if  $A$  then not  $B$ ', and yet (c)  $e_1$  is not improbable given  $e_2$ , and  $e_2$  is not improbable given  $e_1$ .

A Sly Pete argument can be mounted for future indicative conditionals in much the same way as for past indicative condi-

tionals.<sup>8</sup> Although the Sly Pete example is commonly described in the past tense, this is an accidental feature of the example. We imagined your two informants reporting to you after Sly Pete's decision, but we could equally have imagined them to be reporting just beforehand. Informant *A* knows that Sly Pete knows his opponent's hand and says to you, with full justification, 'if Sly Pete calls, he will win'. Informant *B* knows that Sly Pete's hand is the weaker and says, with full justification, 'if Sly Pete calls, he will lose'. And, as before, *A*'s information does not count against *B*'s, and *B*'s information does not count against *A*'s. Thus the case against the natural theory applying to future indicative conditionals is as strong as that against its applying to past indicative conditionals.

#### IV SLY PETE AND PAST SUBJUNCTIVES

An obvious question to ask at this point is whether a Sly Pete argument is possible for past subjunctive conditionals. It had better not be, given our earlier presumption that the natural approach applies to past subjunctives.

I think, however, that we can see that a Sly Pete argument is not possible for past subjunctives. If we take the Sly Pete example and replace the pair of past indicative conditionals by the corresponding pair of past subjunctive conditionals, we get 'if Sly Pete had called, he would have won' and 'if Sly Pete had called, he would have lost'. Now the second of these conditionals is supported by the information that warrants the assertion of the corresponding past indicative (or the corresponding future indicative, if it comes to that). The information that Sly Pete had the weaker hand justifies saying both that if he called, he lost, and that if he had called, he would have lost. However, the first of these past subjunctives — 'if Sly Pete had called, he would have won' — is not warranted by the information about his cheating that warrants the assertion of the corresponding past indicative.

To see this consider the matter from the perspective of informant *A* who is in possession of this information, that is, the information that Sly Pete is cheating by knowing his opponent's hand. We noted before that this information is neutral concerning who has the stronger hand, so let us suppose that *A* gives equal credence to the hypothesis that Sly Pete's hand is stronger, and to the alternative hypothesis that it is weaker than his opponent's. What *A* can be sure about though is that Sly Pete will do the right, in the sense of the most rewarding, thing. That is his return for cheating. But what is the right thing for Sly Pete to do? The

<sup>8</sup> As Gibbard [8] p. 228 notes, though he also notes some complications that space precludes our pursuing.

answer depends on whether or not he has the stronger hand. If he has the stronger hand, the right thing to do is to call; if he has the weaker hand, the right thing to do is to fold. Hence, informant *A* gives a 50 per cent credence to calling being the right action for Sly Pete, and 50 per cent to folding being the right action for Sly Pete. But what makes calling the right action for Sly Pete? Well, he had two options – to call or to fold, and calling is best if it had or would have had the better consequences of the two. But that is the case precisely if had Sly Pete called, he would have won. That is to say, the right credence for informant *A* to give ‘if Sly Pete had called, he would have won’ is the same as the credence he gives to calling being the right thing for Sly Pete to have done, which is 50 per cent. The upshot is that although informant *A*’s information that Pete is cheating makes ‘if Sly Pete called, he won’ very highly assertible, it leaves the credence of the corresponding past subjunctive unchanged at 50 per cent.<sup>9</sup> In brief, informant *A*’s information that Pete is cheating justifies asserting ‘if Sly Pete did not call, then had he called, he would have lost’, but does not justify asserting outright either that had he called he would have lost, or that had he called he would have won.

#### V WHEREIN LIES THE SEMANTIC DIFFERENCE?

Our argument has been that *there* is a semantic difference between past subjunctive conditionals on the one hand, and past and future indicative conditionals on the other: the natural approach applies to the former but not to the latter. But our argument has been silent on wherein lies the difference, for we have not addressed the question of precisely *where* the natural approach fails for the past and future indicative conditionals. There would appear to be three live possibilities. I will not try to adjudicate between them here. I mention them to give a sense of the possibilities opened up by the Sly Pete argument.

You might hold that (Support) fails. Information that (strongly) justifies asserting past and future indicative conditionals does not make them (highly) probably true. The most familiar version of this response holds that past and future indicative conditionals are not truth-valued to start with, and so that there is no question of anything making them probably true. (See e.g. Adams [2], Gibbard [9], and Edgington [5].)

Secondly, you might hold that it is (Conditional non-contradiction) which fails. Even when *A* is consistent, ‘if *A* then *B*’ and ‘if *A* then not-*B*’ can be true together. The most familiar version of this response gives past and future indicative conditionals the truth conditions of the material conditional.

<sup>9</sup> This point has of course played a central role in recent discussions of decision theory, see, e.g., Allan Gibbard and William L. Harper [9].

Finally, you might hold that what is perhaps best described as a presupposition of (Conditional non-contradiction) fails. For you might hold that the truth conditions of past and future indicative conditionals are a function of the context of assertion, that is to say, that the proposition expressed by 'if  $A$  then  $B$ ' varies with that context. More particularly, you might hold that (Conditional non-contradiction) is valid in the following sense: relative to a given context, the proposition expressed by 'if  $A$  then  $B$ ' and the proposition expressed by 'if  $A$  then not- $B$ ' are contraries provided  $A$  is consistent. This is why it is wrong for a given person at a given time to assert both together. However, what may happen is that the proposition expressed by 'if  $A$  then  $B$ ' relative to one context is consistent with the proposition expressed by 'if  $A$  then not- $B$ ' relative to some other context. (See e.g. Stalnaker [14].) And you might hold that exactly this happens in the Sly Pete example: that is to say, you might hold that the proposition that  $A$  expresses by saying 'if Sly Pete called, he won' is consistent with the proposition that  $B$  expresses by saying 'if Sly Pete called, he lost' though it is not consistent with the proposition that  $A$  would have expressed had  $A$  used  $B$ 's words.

(Of course it just might be that the way in which the standard approach fails for future indicative conditionals is different from the way it fails for past indicative conditionals. This is a consequence of the fact that the standard approach has *two* clauses.)

## VI EPISTEMOLOGY AND THE DATA ABOUT AFFIRMATION AND DENIAL

We noted at the very beginning that if we went by affirmation and denial, we would class (1) and (2) together, and so separate them from (3), and I acknowledged an obligation *qua* proponent of classing (2) with (3) semantically, to give an epistemological explanation of the fact that we deny (1), would beforehand deny (2), and affirm (3), and I noted the abstract possibility of doing this. I now want to make this abstract possibility concrete by describing a situation in which we would deny (1) and affirm both (2) and (3), a case, that is, where going by affirmation and denial places (2) and (3) together. It will be apparent that the difference between this situation and our actual one is an epistemological one.

Our justified beliefs about what will happen in the future often depend heavily on our opinions about what the past and present will give rise to. My beliefs about who will win the next election depend heavily on my opinions about the likely consequences of present and past conditions. By contrast, my opinion — indeed, knowledge — about who won last year's election is relatively independent of the more or less speculative views I have about what caused that electoral victory. Similarly, my view beforehand about Lincoln's being killed would normally depend on my view

about how it might happen that he comes to be killed. But we can have exceptions to this norm. I may expect Lincoln to be killed by being killed by Booth working alone, and yet be more confident that he will be killed by someone or other working alone than I am that it will be Booth. Perhaps someone I know to be very knowledgeable about the plot has assured me that Lincoln will be assassinated by someone acting alone, but she remained silent on who it would be. Booth is just my educated guess.

What do I say beforehand? I affirm (2), that if Booth does not kill Lincoln, someone else will. I expect it to be Booth, but be I right or wrong about that part of the story, what I am confident of is that it will be someone or other, and so, if not Booth, then someone else. What do I say afterward when I learn that Lincoln was indeed killed as predicted, that is, by someone working alone, and that it was almost certainly, as I had anticipated, Booth? Clearly I do not affirm that if Booth had not killed Lincoln, someone else would have. Booth acted alone. If he had not killed Lincoln, the plot would have failed and Lincoln would not have been killed. So I deny (1). Finally, provided that although I am confident that it was Booth who killed Lincoln, I am even more confident that Lincoln was killed, I will affirm (3), namely, that if Booth did not kill Lincoln, someone else did.

We have, therefore, concrete reason to hold that the fact that in the normal case we would, going by affirmation and denial, put (2) with (1), and not with (3) has an epistemological explanation not a semantic one, a reason to do with the difference in the epistemic status that normally obtains between opinions about the future and the past. For if we change that status in the appropriate way, going by affirmation and denial puts (2) with (3) not (1).

My case for separating semantically past subjunctive conditionals on the one hand from past indicative conditionals and future indicative conditionals on the other is now before you. Is it theoretically desirable to describe this by saying that the, or anyway an, important line between conditionals is between indicative conditionals and subjunctive conditionals? This depends on what should be said about future subjunctive conditionals, about which we have said nothing, and on the more general question of the viability in English grammar of dividing indicative from subjunctive moods. Here I am content to insist that in respect to one major feature the line has (1) on one side and (2) and (3) on the other.<sup>10</sup>

*Research School of Social Sciences,  
Australian National University, ACT 2601*

<sup>10</sup> I am indebted to too many helpful comments on earlier versions to list, but must mention Lloyd Humberstone, David Lewis, and especially Peter Menzies who gave me invaluable help with (P).

## REFERENCES

- [1] Ernest Adams, 'The Logic of Conditionals', *Inquiry* 8 (1965) 166-97.
- [2] Ernest Adams, *The Logic of Conditionals* (Dordrecht: Reidel, 1975).
- [3] Jonathan Bennett, 'Farewell to the Phlogiston Theory of Conditionals', *Mind* 97 (1988) 509-27.
- [4] V. H. Dudman, 'Parsing "If" – Sentences', *Analysis* 44 (1984) 145-53.
- [5] Dorothy Edgington, 'Do Conditionals have Truth Conditions?', *Critica* 98 (1986) 3-30.
- [6] Brian Ellis, 'A Unified Theory of Conditionals', *Journal of Philosophical Logic* 7 (1978) 107-24.
- [7] Brian Ellis, 'Two Theories of Indicative Conditionals', *Australasian Journal of Philosophy* 62 (1984) 50-66.
- [8] Allan Gibbard, 'Two Recent Theories of Conditionals', in *Ifs*, edited by W. L. Harper *et al.* (Dordrecht: Reidel, 1981), 211-47.
- [9] Allan Gibbard and William L. Harper, 'Counterfactuals and Two Kinds of Expected Utility', reprinted in *Ifs*, edited by W. L. Harper *et al.* (Dordrecht: Reidel, 1981), 153-90.
- [10] Frank Jackson, *Conditionals* (Oxford: Blackwell, 1987).
- [11] David Lewis, *Counterfactuals* (Oxford: Blackwell, 1973).
- [12] Michael Pendelbury, 'The Projection Strategy and the Truth Conditions of Conditional Statements', *Mind* 98 (1989) 179-205.
- [13] Robert Stalnaker, 'A Theory of Conditionals', in *Studies in Logical Theory*, *American Philosophical Quarterly*, monograph series, no. 2, edited by N. Rescher (Oxford: Blackwell, 1968), 98-112.
- [14] Robert Stalnaker, *Inquiry* (Cambridge, Mass.: MIT Press, 1985).