

*Zif Would Have Been If:  
A Suppositional View of Counterfactuals*

David Barnett  
University of Colorado at Boulder

Let us call a statement of the form ‘If *A* was, is, or will be the case, then *C* was, is, or will be the case’ an *indicative conditional*. And let us call a statement of the form ‘If *A* had been, were, or were to be the case, then *C* would have been, would be, or would come to be the case’ a *subjunctive, or counterfactual, conditional*.<sup>1</sup>

On standard views, all conditionals are interpreted as *categorical* statements: to state that, if *A*, then *C*, is to state something, not relative to any hypothesis or supposition, but *categorically*.

Disagreement among the standard views is disagreement over *which* thing is categorically stated by a conditional. For indicatives, H.P. Grice (1967), David Lewis (1976), and Frank Jackson (1979) say that it is the material conditional—that is, the disjunction of the negation of the antecedent with the consequent; Robert Stalnaker (1968) says that it is a predication of a single possible world. For subjunctives, Nelson Goodman (1947) says that it is an entailment from the antecedent, together with laws of nature and particular facts about the actual world, to the consequent;<sup>2</sup> Stalnaker (1968) again says that it is a predication of a single possible world; and Lewis (1973) says that it is an existential generalization over a set of possible worlds.

By contrast, W.V.O. Quine (1950), John Mackie (1973), Michael Dummett (1978), Dorothy Edgington (1995), and I (2006, forthcoming) maintain that conditionals are as their surface form suggests:

---

<sup>1</sup> By this convention, a statement that, if *A* were the case, *C* would be the case, qualifies as counterfactual regardless of whether *A* is actually the case.

<sup>2</sup> More recent views inspired by Goodman’s account include Pollock 1981, Barker 1999, and Hiddleston 2005.

*suppositional* statements.<sup>3</sup> On this view, nothing is categorically stated by a conditional. To state that, if *A*, then *C*, is to state that *C*, *relative to the supposition that A*. What is stated by a conditional is that *C*; what is supposed by it is that *A*. Conditional statements are acts of supposing-cum-stating. This is the suppositional view of conditionals.

I have heard it suggested that the categorical and suppositional views are compatible. On this proposal, a conditional statement can be interpreted as a categorical statement of something, and it can also be interpreted as a statement of something relative to a supposition. To my knowledge, nobody has defended this proposal in print.<sup>4</sup> In any case, it is false: the suppositional and categorical interpretations make incompatible predictions concerning the correct evaluation of conditionals. In my 2006, I draw out some differences in prediction for indicatives; in the present paper, I draw out some differences for subjunctives.

In my 2006 I employ a non-standard method to defend the suppositional view for indicatives. The method is inspired by one of Saul Kripke's ideas:

If someone alleges that a certain linguistic phenomenon in English is a counterexample to a given analysis, consider a hypothetical language which (as much as possible) is like English except that the analysis is *stipulated* to be correct. Imagine such a hypothetical language introduced into a community and spoken by it. *If the phenomenon in question would still arise in a community that spoke such a hypothetical language... then the fact that it arises in English cannot disprove the hypothesis that the analysis is correct for English.* (1977, p. 16)

Kripke's idea is not limited to the evaluation of alleged counterexamples to analyses; it has value for the general process of evaluating analyses. What better way to grasp an analysis, and to assess its consequences, than to examine a community of speakers in which the analysis is stipulated to be correct? The community can be hypothetical, as Kripke suggests, or it can consist of actual speakers who have

---

<sup>3</sup> Views consistent with this idea include Adams 1965, 1966, 1975, Author 2006, Bennett 2003, Dudman 1994, Ryle 1950, and Woods 1997.

<sup>4</sup> To be sure, certain proponents of the categorical view endorse a suppositional method for *evaluating* some conditionals along certain dimensions. But this does not mean that they endorse the suppositional *interpretation* of conditionals. See e.g. Lewis 1973, Stalnaker 1981.

agreed to use a novel expression according to the rules of the stipulation (as with David Kaplan's demonstrative 'dthat' (1978)). If the stipulation breeds phenomena distinct from that typically associated with the analysandum, that is good evidence against the analysis; if it breeds all and only phenomena typically associated with the analysandum, that is good evidence in favor of the analysis.

In my 2006 I stipulate that a sincere and literal utterance of 'zif *A*' is an expression of the supposition that *A*. For indicatives, I argue that 'zif' breeds just the phenomena typically associated with 'if' and that the suppositional view of indicative conditionals is therefore correct. Now I extend my argument to subjunctives. Subjunctives deserve a separate paper because our practices with them are radically different from our practices with indicatives. This is not to suggest that 'if' is ambiguous; 'if *A*' univocally expresses the supposition that *A*. The difference lies in the nature of the supposition expressed. Subjunctive suppositions play a different role in our practices from indicative suppositions. My argument for subjunctives is therefore different from my argument for indicatives.

Hereafter, I limit my attention to subjunctives. The method of stipulating that 'zif *A*' is an expression of the supposition that *A* is especially valuable in this setting. For a significant number of philosophers have become so firmly convinced of the Stalnaker-Lewis nearest possible-worlds analysis, that when asked to evaluate a subjunctive conditional, they first translate the conditional according to the analysis. Only *then* do they evaluate it. Of course, there is little hope of progress if we are conditioned to translate our target expression by a given theory before we evaluate it. By stipulating that 'zif *A*' is an expression of the supposition that *A*, we shall create a context that forces us—at least temporarily—to set aside our theoretical convictions concerning 'if'. In this neutral setting, we can study an analysis that differs fundamentally from the Stalnaker-Lewis analysis. Only after we fully understand the analysis and its consequences will we return to 'if' and attempt to determine, from a theoretically neutral perspective, whether its behavior differs from that of 'zif'.

I begin by reintroducing 'zif' (§1). A number of theoretical questions arise: Do zif-statements have truth-values, probabilities, and modal values? If so, what determines them? Do zif-questions have objectively correct answers? If so, what determines them? These questions concern the evaluation of

suppositional statements and questions—acts of stating and asking things relative to suppositions. As an aid to answering them, I first give a general account of the evaluation of suppositional statements and questions (§2). I then apply this account to the specific case of *zif*-statements and *zif*-questions (§3). Without the benefit of our stipulated definition of ‘*zif*’, a theorist who is in the business of interpreting our use of ‘*zif*’ may encounter a number of *interpretive pitfalls*—features of ‘*zif*’ that are likely to draw the theorist toward an incorrect categorical interpretation, usually of a possible-worlds form. Fortunately for this theorist, our use of ‘*zif*’ is rich with *interpretive clues*—features of ‘*zif*’ that point toward the correct, suppositional, interpretation. I discuss seven pitfalls (§4) and eight clues (§5). Finally, I defend my thesis: that the stipulated role of ‘*zif*’ is the natural role of ‘*if*’, and that the suppositional view of subjunctive conditionals is therefore correct (§6). In other words, if our linguistic community had actually employed ‘*zif*’ according to my stipulation, ‘*zif A*’ would have meant what ‘*if A*’ actually means. For short: *zif* would have been *if*.

## 1. Introducing ‘*zif*’

If you say, “Suppose that Bin Laden had been a pacifist,” I know just what to do. If I reply, “Would the World Trade Center still be standing?” you know just what I have asked. And if you reply, “It probably would be,” I know just what you have stated. We all have an intuitive grasp of what it is to suppose something, and of what it is to state or ask something relative to a supposition. For present purposes, an intuitive grasp will do.

That there is such a phenomenon as stating or asking something relative to a supposition is uncontroversial. Without controversy, then, we need an efficient device to indicate that we are stating or asking something relative to a given supposition. We *could* use the word ‘suppose’ or one of its cognates. Just prior to performing the given act, we could utter (1a), (1b), or (1c):

(1a) I hereby suppose that *A*

(1b) Suppose that *A*

(1c) Supposing that *A*, ...

(1b) and (1c) are more efficient than (1a) because they are less wordy. A minor difference between (1b) and (1c) is that (1b) explicitly invites our audience to suppose that *A*, whereas (1c) merely indicates that we are supposing that *A*. A more significant difference is that only (1c) facilitates scope distinctions. It allows us to indicate, of several claims made in a single utterance, under which suppositions, if any, each is made. Here is an example:

(2) Recent global warming and ozone depletion are results of human activity: supposing that humans had never burned organic matter, recent global warming would probably not have occurred; and, supposing that humans had never developed chlorofluorocarbons, recent ozone depletion would probably not have occurred.

In uttering (2), my first claim is made categorically; my second is made relative to the supposition that humans had never burned organic matter; and my third is made relative to the supposition that humans had never developed chlorofluorocarbons. (1c) is superior to (1a) and (1b), then, because it allows a speaker to selectively indicate, of a number of claims made in a single breath, under which suppositions, if any, each is made.

Still, we can do better. All our candidates so far involve nominalizations—‘that *A*’—and thus involve *reference* to what is supposed. To see that this is inefficient, consider the task of expressing our beliefs. To express the belief that *A*, we could say ‘I have the belief that *A*’ or ‘I believe that *A*’. But this would be inefficient. For all we need to say is ‘*A*’. Our conventions do the rest: the assertoric form of the utterance indicates that we are expressing a belief, and the meaning indicates what we believe. We can in this way communicate our beliefs without reference to them or their contents. Generally speaking, reference burdens the speaker with saying something *about* the referent. In the case of properties, instead of referring to them, thus burdening ourselves with saying something about them, we simply *predicate* them of their bearers. Instead of ‘Fred has the property of being evil’, we simply say ‘Fred is evil’.

It would be useful, then, to have some way to express our suppositions without referring to them. What is needed is a marker to indicate that one who utters a marked clause is supposing its content. The

marker should be lexically compact, unlike (1a). To allow for scope distinctions, it should be embeddable like (1c). And, to avoid reference to what is supposed, it should not nominalize the clause it marks, unlike (1a) – (1c). I propose ‘zif’. Hereafter, by stipulation, a sincere and literal utterance of ‘zif *A*’ shall qualify as an expression of the supposition that *A*. Thus, to ask whether, zif *A*, *C* is to ask whether *C*, relative to the supposition that *A*. And to state that, zif *A*, *C* is to state that *C*, relative to the supposition that *A*.

To get a feel for our new marker, let us rephrase (2) in terms of it:

- (3) Recent global warming and ozone depletion are results of human activity: zif humans had never burned organic matter, recent global warming would probably not have occurred; and, zif humans had never developed chlorofluorocarbons, recent ozone depletion would probably not have occurred.

In terms of speech acts, how do we characterize (3)? A number of acts were performed:

- (3a) I stated, categorically, that recent global warming and ozone depletion are results of human activity.
- (3b) I expressed the supposition that humans had never burned organic matter.
- (3c) Relative to this supposition, I stated that recent global warming would probably not have occurred.
- (3d) I expressed the supposition that humans had never developed chlorofluorocarbons.
- (3e) Relative to this supposition, I stated that recent ozone depletion would probably not have occurred.

And all in a single breath.

Now that I have introduced ‘zif’, a number of theoretical questions arise: Do zif-statements have truth-values, probabilities, and modal values? If so, what determines them? Do zif-questions have objectively correct answers? If so, what determines them? By stipulation, zif-statements and zif-questions are suppositional statements and questions. To answer our questions, then, we need to answer the following questions: Do suppositional statements have truth-values, probabilities, and modal values? If so, what determines them? Do suppositional questions have objectively correct answers? If so, what determines them?

## 2. Evaluating suppositional statements and questions

Statements are acts of stating things. Intuitively, it is the things stated by statements, and not the statements themselves, that are the primary bearers of truth, falsity, probability, necessity and possibility. In practice, we typically evaluate categorical statements by evaluating the things stated by them, not relative to any hypothesis or supposition, but categorically. We typically evaluate suppositional statements by evaluating the things stated by them, not categorically, but relative to the things supposed by them. As theorists, we must pay careful attention to this difference as we investigate the conditions under which various evaluations of suppositional statements are objectively correct.

In the case of categorical statements, the move from evaluating the things stated by statements to evaluating the statements themselves is straightforward. As theorists, we can say that a categorical statement *itself* has a given value just in case a categorical assignment of that value to what is stated by the statement is objectively correct. Of course, a categorical assignment of a value to what is stated by a statement is objectively correct just in case what is stated *has* the value. So, as theorists, we can say that a categorical statement itself has a given value just in case the thing it states has the value. Likewise for categorical questions: a categorical question itself has a given value just in case what it asks has the value.

In the case of suppositional statements, the move from evaluating the things stated by statements to evaluating the statements themselves is less straightforward. For, in practice, the things stated by suppositional statements are stated, not with the intention of being evaluated categorically, but rather with the intention of being evaluated *relative to the given suppositions*. To reflect this practice, we, as theorists, should say that a suppositional statement itself has a given value just in case an assignment of that value to what is stated by the statement, relative to what is supposed by the statement, is objectively correct. Now, generally speaking, an assignment of a value—say, possible, necessary, probable, true, or false—to what is stated by a statement, relative to what is supposed by the statement, is objectively correct, *not* on the condition that what is stated *have* the value assigned, but rather on the condition that what is stated *stand in a corresponding relation to what is supposed*. For example, an assignment of the value *probable*

to the statement *that, supposing hamsters had wings, they would fly* is objectively correct just in case what is stated by the statement—that hamsters would fly—is *made probable* by what is supposed by the statement—that hamsters had wings. Thus, as theorists, we should say that a suppositional statement *itself* has a given value just in case what is stated by the statement stands in a corresponding relation to what is supposed by the statement. Likewise for suppositional questions: a suppositional question itself has a given value just in case what is asked by the question stands in a corresponding relation to what is supposed by the question. For example, the question of how many mountains there would be, supposing there were exactly one more mountain than actual, itself has the value *vague* just in case what is supposed by the question—that there were exactly one more mountain than actual—*entails vagueness* as to what is asked by the question—how many mountains there would be.

So, although we use the same form of words to evaluate categorical and suppositional statements and questions, something very different is going on in the two cases. As theorists, we need to pay careful attention to this difference as we address the following questions: Do suppositional statements have modal values, probabilities, and truth-values? If so, what determines them? Do suppositional questions have objectively correct answers? If so, what determines them? As we answer these questions, we need to keep in mind that our theoretical assignments of values to suppositional statements and questions commit us to the idea that the things stated and asked by them *stand in certain relations* to the suppositions relative to which they are stated and asked. They do *not* commit us to the idea that the things stated and asked *have* the values assigned.

It will be helpful to have an example in hand as we address the preceding questions. Let us say that a *Goldilocks Girl* is a girl who uniquely satisfies the storybook description of Goldilocks. And let us agree that the story of Goldilocks makes no mention of whether Goldilocks likes candy. Now, suppose that there were a Goldilocks Girl. Would she like candy? In other words, supposing there were a Goldilocks Girl, would she like candy? Here is a natural response: “Not necessarily. Of course, it is possible. Indeed, it is even highly probable, for most little girls like candy. It would be silly, however, to

think that there is an objectively correct answer to the question. Strictly speaking, it is neither true nor false that such a girl would like candy.”

In practice, we say that it is not necessary *that, supposing there were a Goldilocks Girl, she would like candy* because what is stated is not *entailed* (i.e., not necessitated) by what is supposed. We say that it is possible *that, supposing there were a Goldilocks Girl, she would like candy* because what is stated is *consistent with* (i.e., left as an open possibility by) what is supposed. To match this practice, in theory we shall say that a suppositional statement itself has the value *necessary/possible* under the following conditions:

|                    |   |
|--------------------|---|
| <i>Necessity</i>   | A suppositional statement is <i>necessary</i> iff what is stated by the statement is entailed by what is supposed by it.    |
| <i>Possibility</i> | A suppositional statement is <i>possible</i> iff what is stated by the statement is consistent with what is supposed by it. |

We may distinguish the logical from the metaphysical modalities by distinguishing logical consistency and entailment from metaphysical consistency and entailment. Remarkably, if we take a claim to be epistemically possible/necessary just in case it is logically *consistent with/entailed by* whatever one is certain of, then, on their logical interpretations, *Necessity* and *Possibility* apply to the epistemic modalities as well. For, in the case of ascriptions of possibility/necessity to suppositional statements, what one is certain of includes the suppositional part of the statement (on the supposition, one should be certain of what is supposed). In practice, this equivalence between epistemic and logical modalities is reflected by our willingness, in the context of a subjunctive supposition, to employ the two modalities interchangeably. In the context of supposing that there were a Goldilocks Girl, we incur the same commitment about the relationship between what is supposed and what is stated by saying any of the following: it is *certain* that she would eat some porridge; she would *definitely* eat some porridge; it is *logically impossible* that she would not eat any porridge; it is *logically necessary* that she would eat some porridge. In practice, there is typically no need to distinguish epistemic from logical modalities when evaluating subjunctive suppositional statements, for ascriptions of either convey the same information; in

practice, we simply say that it is *necessary* that she would eat some porridge and *impossible* that she would not eat any porridge.

We say that it is *probable* that, *supposing there were a Goldilocks Girl, she would like candy* because what is stated is *made probable* by what is supposed. To match this practice, in theory we shall say that a suppositional statement is itself *n% probable* under the following conditions:

*Probability*      A suppositional statement is *n% probable* iff what is stated by the statement is made *n% probable* by what is supposed by it.

Although my argument in this paper does not rest on the question of which theoretical account of *Probability* is correct, three proposals are worth mentioning.

The first is that the probability at issue is ordinary mathematical conditional probability, definable in terms of a ratio of absolute probabilities.<sup>5</sup> This proposal is initially puzzling. For absolute probabilities cannot sensibly be ascribed either to *subjunctive* suppositions or to *subjunctive* claims made relative to them. It does not make sense, for instance, to say that it is *n%* probable that were there a Goldilocks Girl. Nor does it make sense to say, independently of any supposition, that it is *n%* probable that such a girl would like candy. Thus, on the proposal at hand, *Probability* must be interpreted as a claim about a corresponding *indicative* supposition and a corresponding *indicative* statement made relative to it: the probability of a subjunctive claim, *that it would be that C*, relative to a subjunctive supposition, *that it were that A*, is the prior mathematical conditional probability of the *indicative* claim *that it is the case that A*, given the *indicative* claim *that it is the case that C*. I cannot adequately assess this proposal within the scope of the present discussion.

The second proposal is that when we take ourselves to be ascribing probability to subjunctive claims, relative to subjunctive suppositions, what we are really doing is making further subjunctive claims

---

<sup>5</sup> Applied to the case of subjunctive conditionals, this thesis has been held by Adams (1976), Edgington (1995), and Skyrms (1981), among others. For a sample of the sizable literature on the thesis, see Adams 1976, Slote 1978, Skyrms 1981, 1994, Bennett 1984, 2003, Edgington 1995, and Barker 1998.

about probabilities.<sup>6</sup> So, for instance, when I say that it is 95% probable that, supposing there were a Goldilocks Girl, she would like candy, what I am really saying (or trying to say) is that, supposing there were a Goldilocks Girl, it *would be* 95% probable that she would like candy. This proposal is implausible, for at least two reasons.

First, it fails to capture the intention behind my ascription of probability. Suppose that there were a Goldilocks Girl. It is possible that such a girl would like candy, and it is possible that such a girl would not like candy. Given the facts about children and their tendency to like candy, the former possibility *is* more likely than the latter: that such a girl would like candy *is* more likely than that such a girl would not like candy. When I say that it *is* 95% probable that such a girl would like candy, I do not intend to be saying something about how probable some claim *would be*, supposing that there were a Goldilocks Girl. Rather, I intend to be saying something about how probable it *is* that a given possibility *would be* realized. The first reason the second proposal is implausible, then, is that it fails to capture the intention behind ascriptions of probability to subjunctive suppositional statements.

To see the second reason, note that corresponding to my ascription of probability is an ascription of confidence: I am 95% confident that, supposing there were a Goldilocks Girl, she would like candy. By adapting the suggestion to apply to degrees of confidence, we arrive at the view that my ascription of confidence is *really* a subjunctive claim about how confident I *would be* that such a girl would like candy. But this is absurd. Surely my 95% confidence that such a girl would like candy does not commit me to the view that, supposing there were a Goldilocks Girl, I *would be* 95% confident that such a girl would like candy. For it does not even commit me to the view that I would exist, supposing there were a Goldilocks Girl. To emphasize the point, suppose that there were no conscious beings. Presumably, you are 100% confident that you would not exist. This confidence does not commit you to the absurd claim that, supposing there were no conscious beings, you *would be* 100% confident that you would not exist. So the suggestion about ascriptions of probability cannot be adapted to the case of degrees of confidence. Once it is admitted that our ascriptions of degrees of confidence should be taken at face value, as ascriptions of

---

<sup>6</sup> Bennett (2003, p. 252) endorses this idea as applied to subjunctive conditionals.

confidence to what is stated, relative to what is supposed, it is hardly plausible that our corresponding ascriptions of probability should receive a fundamentally different treatment, on which they are reinterpreted as having a logical form entirely different from their surface form, and on which they are really not even ascriptions of probability at all, but rather bare subjunctive claims about hypothetical ascriptions of probability. This is the second reason to doubt the second proposal.

On the third proposal, which I favor, our ascriptions of probability are taken at face value, as indicative ascriptions of probability, to subjunctive claims, relative to subjunctive suppositions. Call a probability that such and such would be the case, relative to the supposition that so and so were the case, a *Conditional Counterfactual Probability*—a *CCP* for short. What facts determine CCPs? Conversely, what do CCPs measure?

CCP's appear to measure the *stability* of features and connections in the world. Suppose for illustration that a large number of children have been surveyed and that 95% of them like candy. The question arises whether this statistic reflects a relatively stable connection between being a child and liking candy, or whether it is purely accidental. First suppose that it reflects a stable connection. Then we are inclined to say that it is 95% probable that the next child surveyed will like candy and, supposing that there had been one additional child in the world, it is 95% probable that that child would have liked candy. These ascriptions of probability reflect our confidence that the connection between being a child and liking candy is relatively stable. Suppose now that the connection is purely accidental. Then there is no longer a clear motivation for ascribing 95% probability to either claim.

Should there in fact be a stable connection between being a child and liking candy, the connection is not basic, but rather depends on a wide variety of other more basic connections, connections for instance between being a human and having a given physiological make-up, being a piece of candy and having a given chemical make-up, and being a system of given molecular sort and being a system that evolves in a given sort of way. The relatively stable connections give way to ones that are more stable, more general, and more basic, until ultimately we reach the brute stabilities, including the fundamental laws of nature. That some stable connections depend on others is reflected by certain CCP's. For instance,

supposing that most children were averse to candy, it is very likely that the gravitational law would be the same. But, supposing that the gravitational law were different, it is very unlikely that children would even exist.

This characterization of CCP's might help to explain the appeal of the first proposal that (i) *the probability of a subjunctive suppositional statement* is identical to (ii) *the prior mathematical conditional probability of the corresponding indicative statement, given the corresponding indicative supposition*.

For, by this characterization, (i) measures the stability of features and connections in the world. And there is some initial appeal to the idea that (ii) measures the very same thing: just as stable features and connections are likely *to have obtained*, all things equal, they are likely *to continue to obtain*, all things equal.<sup>7</sup>

I have considered three proposals for characterizing the probability involved in *Probability*. My argument does not rest on which, if any, of the three is correct. We can comfortably endorse *Probability* without giving any theoretical account of it.

Returning to our main line of discussion, we say that *there is no answer* to the question of *whether, supposing there were a Goldilocks Girl, she would like candy* because what is supposed *entails no answer* to what is asked. To match this practice, in theory we shall say that a suppositional question itself has the value *factual* under the following conditions:

*Factuality*      A suppositional question is *factual* iff what is supposed by the question entails an answer to what is asked by it.

In conformity with *Factuality*, we say that it is *neither true nor false* that, *supposing there were a Goldilocks Girl, she would like candy* because what is stated is neither *made true by*—that is, *entailed by*—what is supposed nor *made false by*—that is, *inconsistent with*—what is supposed. To match this

---

<sup>7</sup> Indeed, Skyrms (1981) suggests that problems facing early versions of the “prior probability” view of subjunctive conditionals (the view corresponding to the first proposal above) can be solved by replacing subjective interpretations of probability with an objective “propensity” interpretation—one that is more in line with the idea that the probabilities measure the stability of features in the world.

practice, in theory we shall say that a suppositional statement is itself *true/false* under the following conditions:

|                |  |
|----------------|--|
| <i>Truth</i>   | A suppositional statement is <i>true</i> iff what is supposed by the statement entails what is stated by it.               |
| <i>Falsity</i> | A suppositional statement is <i>false</i> iff what is supposed by the statement is inconsistent with what is stated by it. |

Initially, these truth conditions might appear too demanding. For we often maintain a high degree of confidence in a suppositional statement even when it is obvious to us that what is supposed by the statement does not entail what is stated by it. We might, for instance, maintain a high degree of confidence that, supposing that there were a Goldilocks Girl, she would like candy, despite having zero degree of confidence that the supposition *that there were a Goldilocks Girl* entails *that such a girl would like candy*. Should we conclude, on this basis, that there are counterexamples to *Truth*? Only if confidence in a suppositional statement rationally requires confidence, of the statement, that it is true. But it does not. In practice, confidence in a suppositional statement is confidence in what is stated by the statement, relative to what is supposed. This conditional confidence commits us to the idea that what is supposed *makes probable* what is stated. By contrast, an ascription of truth to a suppositional statement is an ascription of truth to what is stated by the statement, relative to what is supposed. This conditional ascription of truth commits us to the idea that what is supposed *makes true* what is stated. (Keep in mind that we are limiting our attention to subjunctives.) Because what is supposed might make what is stated probable without making it true—that is, without entailing it—confidence in a suppositional statement need not rationally require confidence, of the statement, that it is true. For instance, on the supposition that there were a Goldilocks Girl, it need not be irrational to be confident that such a girl would like candy, while maintaining that there is no objectively correct answer to the question of whether such a girl would like candy, and that it is therefore neither true nor false that such a girl would like candy. So, the fact that we often have high degrees of confidence in suppositional statements, even when we when have

zero degree of confidence that what is supposed by the statements entails what is stated, does not show that the proposed truth conditions are too demanding.<sup>8</sup>

*Truth* is perhaps worth further motivating with an additional example. Imagine that Ella and Ed are sharing a glass of champagne on the balcony of their 34<sup>th</sup> story New York City apartment. As they lean over the edge of the balcony to gaze at the street below, the glass slips out of Ed's hand. Fortunately, Ella snatches it out of the air without spilling a drop.

Ed: That was a close one.

Ella: Indeed. Suppose the glass had fallen!

Ed: It definitely would have shattered.

Ella: Well, I hate to be a stickler, but I don't think it's right to say that it *definitely* would have shattered. For, as unlikely as it sounds, a perfect gust of wind could have brought the glass to a gentle landing. Or a pedestrian with the perfect hairdo for bringing the glass to a gentle landing could have been standing directly beneath us. To be sure, it is *highly unlikely* that either of these scenarios would have obtained. But either *could* have obtained. And so, strictly speaking, it is incorrect to say that the glass *definitely* would have shattered.

Ed: That's ridiculous. In fact, there was no such gust of wind, no such pedestrian, and no other mechanism in place to prevent the glass from shattering. So it *is* correct to say that the glass *definitely* would have shattered.

Ella: From the fact that certain conditions did not obtain, it does not follow that they *would not have obtained*, supposing that things had been different. After all, the glass did not fall, but it definitely *would have* fallen, on our supposition. Likewise, there was in fact no such gust of wind, but there *could have been* such a gust, supposing that things had been different in a way that is consistent with such a gust. For emphasis, just think of a couple of the ways that the glass could have fallen. It could have fallen due to a subtle difference in the initial conditions of the universe, say, one that led to your reactions being a bit slower than they actually were. This difference could also have led to the existence of a perfect gust of wind. Another way that the glass could have fallen is for there to have been a subtle difference in the laws of nature, say, one that led to the glass's accelerating slightly faster than it actually did. This difference, too, could have led to the existence of a perfect gust of wind. So, although the facts about what actually happened might make it *highly likely* that the glass would have shattered, none *guarantees* that it would have shattered.

---

<sup>8</sup> How can we be said to *know* that, supposing that *A* had been the case, *C* would have been the case, when what is supposed does not entail what is stated? Answer: we cannot. If there is no fact of the matter whether *C* would have been the case, then there is nothing to know. Still, we can be justifiably confident that *C* would have been the case. For the supposition *that A had been case* may make it probable *that C would have been the case*.

- Ed: Fine. I suppose that nothing *guarantees* that it would have shattered, and so—strictly speaking—it is not correct to say that it *definitely* would have shattered. Still, I bet that it would have shattered.
- Ella: Just to humor you, I'll take you up on that bet: I bet you \$100 that it wouldn't have shattered.
- Ed: Are you crazy? We are on the 34<sup>th</sup> story! We already agree that it is extremely likely that it would have shattered, right? So likely, in fact, that we might as well agree that it would have shattered. Hence, we might as well agree that I've won the bet.
- Ella: Unfortunately for you, the fact that it is *extremely likely* that it would have shattered does not settle who has won the bet. It settles the question of whether it is extremely likely that the glass would have shattered. But we did not bet on this question. We bet on *whether the glass would have shattered*. For you to win *this* bet, some fact must make it the case—that is, guarantee—not merely that it is *extremely likely* that the glass would have shattered, but *that the glass would have shattered*. And we have already agreed that no fact *guarantees* that the glass would have shattered. So your bet is not a winning bet. Of course, neither is mine. But, as I said, it was only to humor you that I made the bet in the first place.
- Ed: Point well taken. I suppose a similar line of thought shows that it is not *true* that the glass would have shattered. For surely it is true that the glass would have shattered only if my bet is a winning bet. So not only is it incorrect to say that the glass *definitely* would have shattered, it is also incorrect to say that it is *true* that the glass would have shattered, despite the fact—on which we have agreed all along—that it is *highly likely* that the glass would have shattered.
- Ella: Bingo. Given the facts surrounding the actual scenario, our supposition *makes likely* that the glass would have shattered, but it does not guarantee it—that is, it does not *make it true*.

Before proceeding, it is perhaps worth emphasizing a certain danger in engaging in the theoretical practice of assigning truth conditions to suppositional statements in the first place. Even upon recognizing that *Truth* is an account of the truth conditions of suppositional statements themselves, and not of the things they state, one might be tempted to infer from *Truth* that suppositional statements can be treated as categorical statements of things that are true under the conditions given by *Truth*. After all, once absolute truth conditions are assigned—to suppositional statements themselves, what prevents us from treating these statements as categorical statements of things with the very same truth conditions? More generally, what prevents us from treating suppositional statements *both* as statements of things relative to suppositions *and* as categorical statements of things?

One way to see the answer to this question is to consider various *combinations* of assignments of values to what is stated by a suppositional statement, relative to what is supposed. These combinations commit us to various combinations of relations between what is supposed and what is stated. On a categorical interpretation, they also commit us to the existence of a single content which itself *has* the combination of values. But in §5 we will find combinations of values such that (i) an assignment of them to a suppositional statement commits us only to unproblematic combinations of relations between what is supposed and what is stated; even though (ii) no single content could itself have the combination of values. Thus, for certain combinations of values, it may be objectively correct to assign the combination to a suppositional statement, even though it could not be objectively correct to assign it to a categorical statement.<sup>9</sup> This is one reason why suppositional statements cannot be interpreted categorically.

We began this section by asking: Do suppositional statements have modal-values, probabilities, and truth-values? If so, what determines them? Do suppositional questions have objectively correct answers? If so, what determines them? To these questions I gave the following answers: *Necessity*, *Possibility*, *Probability*, *Truth*, *Falsity*, and *Factuality*.

### **3. Evaluating zif-questions and zif-statements**

Zif there were a Goldilocks Girl, would she like candy? Not necessarily. Of course, it is possible. Indeed, it is even highly probable, for most little girls like candy. It would be silly, however, to think that there is an objectively correct answer to the question. Strictly speaking, it is neither true nor false that, zif there were a Goldilocks Girl, she would like candy.

By stipulation, zif-questions and zif-statements are suppositional questions and statements. Thus, our account of the evaluation of suppositional questions and statements applies straightforwardly to the case of zif-questions and zif-statements:

---

<sup>9</sup> In §5 the point is formulated in terms of a difference in the laws governing the rational evaluation of suppositional versus categorical statements.

|                        |   |
|------------------------|---|
| <i>Zif Necessity</i>   | A zif-statement is <i>necessary</i> iff what is stated by the statement is entailed by what is supposed by it.                  |
| <i>Zif Possibility</i> | A zif-statement is <i>possible</i> iff what is stated by the statement is consistent with what is supposed by it.               |
| <i>Zif Probability</i> | A zif-statement is <i>n% probable</i> iff what is stated by the statement is made <i>n%</i> probable by what is supposed by it. |
| <i>Zif Factuality</i>  | A zif-question is <i>factual</i> iff what is supposed by the question entails an answer to what is asked by it.                 |
| <i>Zif Truth</i>       | A zif-statement is <i>true</i> iff what is supposed by the statement entails what is stated by it.                              |
| <i>Zif Falsity</i>     | A zif-statement is <i>false</i> iff what is supposed by the statement is inconsistent with what is stated by it.                |

#### 4. Interpretive pitfalls

Our new device presents both hazards and opportunities for those I call *outsiders*—theorists who, without the benefit of our stipulated definition, are in the business of interpreting our use of ‘zif’. To make their business realistic, we shall cooperate in their investigation only as competent users of ‘zif’, and not as theorists with explicit knowledge of its definition.

The greatest challenge facing the outsider is to recognize that zif-clauses express suppositions and that speech acts that begin with the utterance of a zif-clause are therefore suppositional in nature—they are *not* acts of relating *categorically* to their contents. Zif-statements are not acts of stating categorically what they state, and zif-questions are not acts of asking categorically what they ask. In this section, we consider features of ‘zif’ that are likely to draw the outsider toward a categorical interpretation, usually of a possible-worlds form; these are the *interpretative pitfalls of zif*. In the next section, we consider features of ‘zif’ that point toward its correct interpretation; these are the *interpretive clues of zif*.

Seven pitfalls are worth highlighting.

*Pitfall #1.* In practice, we never speak of, or evaluate, the content of a suppositional statement or question—that is, what is stated or asked by the statement or question—apart from its suppositional context—that is, what is supposed by the statement or question. This is because the identity and

significance of a suppositional statement/question depends not just on (i) the *relation* borne by the agent (stating/asking) and (ii) the *content* to which the agent bears the relation (what is stated/asked by the agent), but also on (iii) the suppositional *context* in which the relation is borne (what is supposed by the agent). In practice, it would be pointless to consider (i) and (ii) apart from (iii), for neither the act in question nor its significance could be identified.

In the case of indicatives, it would be *misleading* to report (i) and (ii) apart from (iii), for one's audience would falsely assume that the relevant act had no suppositional context. If you witness Emma state that, supposing that God exists, atheism is false, and later I ask you what Emma stated, you are obliged to tell me not just that Emma stated that atheism is false—this would be misleading—but also that she stated this *on the supposition* that God exists. I ask, “What did Emma state?” You answer, “That, zif God exists, atheism is false,” thus giving me all the information I need to identify the act and its significance.

In the case of subjunctives, it is *impossible* to report (i) and (ii) apart from (iii), for the content of a subjunctive suppositional statement depends essentially on the suppositional context. If you witness Emma state that, supposing that hamsters had wings, hamsters would fly, and later I ask you what Emma stated, you are obliged to tell me not just that Emma stated that hamsters would fly—this would not be enough information to convey even the content of her statement—but also that she stated this *on the supposition* that hamsters had wings. I ask, “What did Emma state?” You answer, “That, zif hamsters had wings, they would fly,” thus giving me all the information I need to identify the act and its significance.

Because in practice we never speak of the content of a suppositional statement or question apart from its suppositional context, semantic information that in fact determines suppositional context might, to the outsider, appear to play a role solely in determining what is stated or asked by the act. This appearance might blind the outsider to the fact that zif-clauses express suppositions. As a result, the outsider might conclude that a statement that, zif *A* were the case, *C* would be the case, is a categorical statement of something whose identity is a function of the meanings of ‘*A*’ and ‘*C*’.

*Pitfall #2.* The second pitfall is that a high degree of confidence in a zif-statement is rationally consistent with a low degree of confidence that the statement is true. This is a remarkable pitfall, for typically the outsider will theorize under the assumption that our high degrees of confidence in zif-statements indicate that we are confident that such statements are true. On this basis, he is likely to conclude that zif-statements are categorical statements of things whose truth does not generally require the entailment of the consequents by the antecedents; for we will tend to be confident in a variety of zif-statements whose antecedents do not entail their consequents.

*Pitfall #3.* Early in his investigation, the outsider is likely to ask us what zif-statements are about. He might for instance ask what statement (4) is about:

(4) that, zif hamsters had wings, they would fly.

A natural answer is that (4) is about a hypothetical scenario in which hamsters have wings; in that scenario, hamsters would fly. One way to state something, relative to a given supposition, is to state something *about a corresponding hypothetical scenario*. Once we suppose that such and such were the case, we have at hand a corresponding hypothetical scenario, one in which such and such is the case. Talk of how things would be *in that scenario* is equivalent to talk of how things would be *on the given supposition*. Because hypothetical-scenario talk is equivalent to suppositional talk, and because suppositional talk cannot be interpreted categorically, hypothetical-scenario talk cannot be interpreted categorically. (More in defense of this claim below.) Yet, on the surface, hypothetical-scenario talk looks very much like categorical talk. Just as we might describe, categorically, how things are *at a given time* or *in a given space*, it might appear that we can describe, categorically, how things would be *in a given hypothetical scenario*. This appearance, together with our disposition to say that zif-statements are about hypothetical scenarios, might lead the outsider to conclude that zif-statements are categorical statements about hypothetical scenarios.

*Pitfall #4.* Hypothetical scenarios come in two varieties: possible and impossible. Possible scenarios are those that could be realized; impossible scenarios are those that could not. We say, of

possible scenarios, that they depict *possible situations*. Zif-statements are typically about possible, rather than impossible, scenarios; in turn, they are typically about possible situations. (4), for instance, is about a possible situation in which hamsters have wings; in that situation, hamsters would fly. So, when asked what (4) is about, there are at least two good answers: (i) it is about a *hypothetical scenario* in which hamsters have wings, or (ii) it is about a *possible situation* in which hamsters have wings. Because talk of how things would be *in a possible situation* is equivalent to talk of how things would be *in a hypothetical scenario that depicts the situation*, and because the latter cannot be interpreted categorically, the former cannot be interpreted categorically. (More in defense of this claim below.) Still, as with hypothetical-scenario talk, possible-situation talk looks superficially like categorical talk. And so the outsider might conclude that zif-statements are categorical statements about possible situations.<sup>10</sup>

*Pitfall #5.* Possible hypothetical scenarios come in two varieties: uniquely realizable and multiply realizable. A uniquely realizable scenario depicts a perfectly specific and complete situation; in philosopher's jargon, it depicts a *possible world*. A multiply realizable scenario depicts a general and incomplete kind of situation, one that is itself realizable by numerous possible worlds. Because (i) hypothetical-scenario talk looks superficially like categorical talk; (ii) zif-statements are typically about possible hypothetical scenarios; and (iii) a possible hypothetical scenario is realizable either by a unique possible world or by a set of possible worlds, the outsider might conclude that a zif-statement is a categorical statement about a possible world or set of possible worlds.<sup>11</sup>

*Pitfall #6.* Often, the author of a zif-statement will intend to make her statement in an enriched suppositional context, one that includes not just the supposition literally expressed by her zif-clause, but also the supposition *that everything else were as similar as possible to actual* (at least up to the specified point of subjunctive divergence from actuality, should there be one). Call this extra condition a *nearness condition*. For illustration, imagine that two officers, Smith and Jones, accompany fifty soldiers to the edge of a minefield. Smith orders the soldiers across the field. The field is so densely mined that only one

---

<sup>10</sup> For this interpretation of subjunctive conditionals, see Lycan 2001.

<sup>11</sup> For this interpretation of subjunctive conditionals, see Stalnaker 1968 and Lewis 1973.

of the fifty makes it more than halfway. Unfortunately, just as this soldier reaches the end of the field, she steps on a mine and is blown to smithereens.

Smith: Zif she hadn't stepped on that mine, she would have made it across.

Jones: I doubt it. For suppose that she hadn't stepped on that mine. We must ask ourselves: what is the mostly likely way for this to have come about? Perhaps the initial conditions of the universe had been different; in which case it is highly unlikely that she, or this minefield, would ever have existed. Perhaps the initial conditions had been the same, but the laws had been different; in which case it is again highly unlikely that she, or this minefield, would ever have existed. Perhaps the initial conditions and the laws had been the same, but—perhaps owing to a miracle or to indeterministic laws—the world had nevertheless evolved differently at some point; in which case it is again highly unlikely that she, or this minefield, would ever have existed. Perhaps the world had evolved just as it actually did up until the point at which you ordered the soldiers across the field; in which case the most likely way for her to have not stepped on that mine is for her to have stepped on an earlier mine—after all, she was lucky to get as far as she actually did; in which case, once again, she would *not* have made it across the field. Do you get the picture? Zif she hadn't stepped on that mine, it is extremely likely that it would have been owing to circumstances in which she would *not* have made it across the field.

Smith: You are extremely uncharitable. Was it not obvious from our context that what I *meant* was that, zif she hadn't stepped on that mine *and things had been as similar as possible to actual, up to that point*, she would have made it across?

Jones: Well, in *that* case, she probably *would* have made it across. From now on, please *say* exactly what you *mean*.

Because a nearness condition is often implicit in the context of a zif-statement, authors of zif-statements will often be disposed to say that their statements are about hypothetical scenarios in which the antecedents hold *and all else is as similar as possible to actual* (hereafter I suppress the proviso about the point of subjunctive divergence). Combined with the preceding pitfalls, this might lead the outsider to conclude that zif-statements are categorical statements about the *nearest* possible world or set of worlds in which the antecedent holds.<sup>12</sup>

*Pitfall #7.* This pitfall might lead the outsider to the same conclusion. It centers on what I call *subjunctive freewill suppositions*—suppositions on which agents were to freely act differently from how they in fact freely act. These suppositions are of great practical importance. By reasoning with them, we

---

<sup>12</sup> For this interpretation of subjunctive conditionals, see Stalnaker 1968 and Lewis 1973.

learn how to better control our futures. For illustration, imagine that Joe suffers a serious injury from a car accident. He considers what would have happened zif he had worn a seatbelt. On reflection, he becomes confident that he would not have suffered the injury. Through this exercise in subjunctive suppositional reasoning, he becomes confident that there is a stable connection between wearing a seatbelt and driving without injury. Joe then chooses to wear his seatbelt in the future. Because subjunctive freewill suppositions play an important role in how we reason about stable features and connections in the world, and thus how we learn to control our futures, the outsider is likely to encounter zif-statements involving them. To begin to see the unique pitfall that awaits him, consider a continuation of the preceding exchange between Smith and Jones:

Smith: I *could* have explicitly added a nearness condition to the suppositional part of my statement. Alternatively, however, I could have made a different sort of adjustment to my initial statement. What I could have said is that, zif the soldier had *freely chosen* to step just to the left of where she actually *freely chose* to step (when she stepped on the mine), she probably would have made it across.

Jones: Very clever. Assuming that her choice to step where she did was indeed a *free* choice, then your adjustment obviates the need to include a nearness condition in your supposition. For to say that her choice was free is to say that it was free of outside influence, which suggests that there were no stable contingent connections between it and any events that preceded it. So, on the supposition that it had been different, there is no pressure to explain this difference in terms of a difference in events leading up to it. Moreover, there are two reasons to think the events leading up to it would have been just as they actually were. First, the soldier had to have been presented with the opportunity to choose to step just to the left of where she actually stepped, and the actual series of events leading up to her choice would have presented her with just this opportunity. Second, it is plausible that the events leading up to her choice did not, as a matter of fact, happen by pure accident, but were rather underlain by certain stable features and connections in our world. So, zif the soldier had *freely chosen* to step just to the left of where she actually *freely chose* to step, the events leading up to this choice would probably have been *just as they actually were*, for there is no reason to think they would have been different, and there is some reason to think they would have been the same. To be sure, there is a chance that they would have been different, for their having been the same is not entailed by our supposition. Still, without loading a nearness condition into the supposition, we can be confident—though not certain—that a nearness condition would have obtained.

More generally, given a conception of freewill on which a free act is one that is free of outside influence, speakers can be confident—though not certain—that, zif an agent had freely chosen to do something different from what she actually freely chose, things leading up to her choice would have been

just as they actually were. Because such a conception of freewill is common (at least outside professional philosophy), and because zif-statements with freewill suppositions are of great practical value, the outsider is likely to encounter this sort of confidence in his investigation. If he interprets this confidence as confidence that the relevant zif-statements are true, and if he generalizes on this phenomenon to all zif-statements, he is likely to conclude, once more, that zif-statements are categorical statements about the *nearest* possible world or set of worlds in which the antecedent holds.

The outsider who arrives at this conclusion is wrong on three counts. First, zif-statements do not generally have a nearness condition built into their semantics; hence, not all zif-statements are about the *nearest* possible world or set of worlds in which the antecedent holds. Second, not all zif-statements are about *possible* hypothetical scenarios; hence, not all zif-statements are about possible worlds. Third, and most importantly, zif-statements are not categorical statements.

## 5. Interpretive clues

Fortunately for the outsider, our use of ‘zif’ is rich with features pointing toward its correct interpretation. These are the interpretive clues of ‘zif’. In this section I highlight eight clues. The first reveals that zif-statements do not generally have a nearness condition built into their semantics. The second reveals that not all zif-statements are about *possible* hypothetical scenarios. Each of the last six reveals that zif-statements are not categorical: essentially, they state things *relative to suppositions*.

*Clue #1.* To test the hypothesis that zif-statements are by nature about the nearest possible world or set of worlds in which the antecedent holds, the outsider might investigate whether explicitly adding a nearness-condition to the antecedent of a zif-statement has any effect on our evaluation of the statement. On the nearest-world hypothesis, it should not. The outsider might for instance ask us to compare (5a) with (5b):

- (5a) Zif hamsters had wings, everything else would be as similar as possible to actual.
- (5b) Zif hamsters had wings and everything else were as similar as possible to actual, everything else would be as similar as possible to actual.

As competent users of ‘zif’, we evaluate (5a) and (5b) differently.

First, we assign different modal values to the two. It is not necessary that, zif hamsters had wings, everything else would be as similar as possible to actual. Zif hamsters had wings, everything else could have been very different from actual. Clearly, (5a) is not necessary. By contrast, (5b) *is* necessary: necessarily, zif hamsters had wings *and everything else were as similar as possible to actual*, then everything else would be as similar as possible to actual. So, we assign different modal values to (5a) and (5b).

Second, we assign different probabilities. First consider (5a). Although it may be highly likely that, zif hamsters had wings, everything else would be as similar as possible to actual, it is not 100% likely. Zif hamsters had wings, everything else might have been very different from actual. In other words, there is some chance that, zif hamsters had wings, everything else would be very different from actual. So (5a) is less than 100% probable. By contrast, (5b) is 100% probable. So, we assign different probabilities to (5a) and (5b).

The first clue for the outsider, then, is that competent users of ‘zif’ evaluate (5a) and (5b) differently. This is evidence that zif-statements do not generally have a nearness condition built into their semantics, which is evidence that zif-statements are not by nature about the *nearest* possible world or set of worlds in which the antecedent holds. (To be sure, on the suppositions expressed by a significant class of zif-statements, it might be *highly probable* that things would have been as similar as possible to actual; but this is consistent with the fact that such statements have no nearness condition built into their semantics.)

*Clue #2.* Competent users of ‘zif’ will attest to the fact that some meaningful zif-statements, including the following, are about *impossible* scenarios:

- (6a) Zif the truths of fundamental physics were discoverable by a priori conceptual analysis, particle accelerators would be superfluous.

The fact that we deem (6a) meaningful, and yet take it to be about an impossible scenario, is evidence that zif-statements are not by nature about possible scenarios. Still, given our high degree of confidence in (6a), the outsider might misinterpret this particular evidence as support for the hypothesis that, while zif-statements by nature *purport* to be about possible scenarios, they sometimes fail, and when they do fail, they are vacuously true.<sup>13</sup> For the full clue, then, the outsider needs to consider our reactions to a variety of zif-statements with impossible antecedents. Eventually he will see that we often ascribe low degrees of confidence to such statements, as we do to (6b):

(6b) Zif having wings were to entail being a bird, having wings would not entail being a bird.

The second clue, then, is that we judge some zif-statements to be about impossible scenarios, and our confidence in such statements is sometimes low and sometimes high. This does not comport with the hypothesis that zif-statements with impossible antecedents are vacuously true (or vacuously false). In light of this clue, the outsider is well advised to reject the hypothesis that zif-statements by nature purport to be about *possible* scenarios.<sup>14</sup>

The last six clues reveal that zif-statements are suppositional, rather than categorical, statements. The first five center on a general difference between categorical and suppositional statements. When we ascribe a value to a categorical statement, we commit to the idea that what is stated by the statement *has* the value. But when we ascribe the very same value to a suppositional statement, we commit to the idea that what is stated *stands in a certain relation* to what is supposed. This difference gives rise to differences in the laws governing the rational evaluation of categorical and suppositional statements. For instance, there are pairs of values such that it may be rationally permissible to ascribe the pair to a suppositional, but not to a categorical, statement. Differences of this sort give rise to a plethora of clues for the outsider. Below I limit my attention to just four such clues, #3 – #6; however, it should become clear to the reader that these four are members of a larger family.

---

<sup>13</sup> For this interpretation of subjunctive conditionals, see Stalnaker 1968 and Lewis 1973.

<sup>14</sup> Applied to the case of subjunctive conditionals, a proponent of the possible-worlds theory might employ *impossible* worlds to accommodate the relevant phenomena; see Nolan 1997 and Kment 2006.

*Clue #3.* The following principle governs the rational evaluation of categorical, but not suppositional, statements:

- (7a) However confident one is that *S*, one should be at least as confident that there is an answer to the question of whether *S*.

To see that (7a) holds for categorical statements, consider an arbitrary categorical statement, say, that the sky is blue. However confident you are that the sky is blue, you should be equally confident that the answer to the question of whether the sky is blue is that the sky is blue (at least on consideration of the matter). Thus, however confident you are that the sky is blue, you should be at least as confident that there is an answer to the question of whether the sky is blue. Clearly (7a) holds for categorical statements.

(7a) does not hold for suppositional statements. Relative to a subjunctive supposition *that A were the case*, it can be rational to be (i) highly confident *that C would be the case* and yet (ii) certain that there is no answer to the question *whether C would be the case*. (i) commits one to the view that the supposition makes it highly probable that *C* would be the case. (ii) commits one to the view that the supposition entails no answer to the question whether *C* would be the case. There is nothing problematic about a supposition which makes something highly probable without entailing it or its negation, and (i) and (ii) have no further jointly problematic commitments. So it is easy to see why (7a) does not govern the rational evaluation of subjunctive suppositional statements.

Competent users of ‘zif’ will attest to the fact that (7a) does not govern the rational evaluation of zif-statements. For instance, we will attest to the fact that a high degree of confidence that, zif there were a Goldilocks Girl, she would like candy, does not rationally require a high degree of confidence that there is an answer to the question whether, zif there were a Goldilocks Girl, she would like candy. To be sure, we will endorse a similar principle:

- (7b) However confident one is that, zif *A* were the case, *C* would be the case, one should be at least as confident that, zif *A* were the case, there *would be* a question as to *whether C was the case*, which *would* have an answer.

For instance, however confident one is that, zif there were a Goldilocks Girl, she would like candy, one should be at least as confident that, zif there were a Goldilocks Girl, there would be an answer to the question of whether she liked candy. Our confidence that things would be a certain way, on the supposition that *A* were the case, need not rationally translate into confidence that things *are in fact* a certain way with respect to some zif-question, but rather into confidence that things *would be* a certain way with respect to an ordinary categorical question that *would* exist, zif *A* were the case. The outsider who distinguishes (7a) from (7b) will be rewarded with the third clue: competent users of ‘zif’ will attest to the fact that the rational evaluation of zif-statements is governed by (7b) rather than (7a). This reveals that zif-statements are suppositional, rather than categorical, statements.<sup>15</sup>

*Clue #4.* The following principle governs the rational evaluation of categorical, but not suppositional, statements:

- (8) On the supposition that there is no answer to the question of whether *a* is *F*, one should have zero confidence that *a* might be *F*.

To see that (8) holds for categorical questions, consider an arbitrary categorical question, say, whether the sky is blue. Suppose that there is no answer to this question. Then one cannot rationally have any confidence that the sky *might* be blue, for any such confidence rationally requires at least as much confidence that there might be an answer to our question—namely, that the sky *is* blue. And we are supposing that there is no answer to this question. Generalizing, (8) holds for all categorical statements.

To see that (8) does not hold for suppositional statements, suppose that there were a Goldilocks Girl. One of two possibilities would be realized: either she would like candy or she would not. Now suppose that there is no answer to the question whether such a girl would like candy, and thus no answer to the question which of these two possibilities would be realized. Still, each *is* a possibility: such a girl *might* like candy, and such a girl *might not* like candy. Another way of putting the point is that our present epistemic situation, constituted by the supposition that there were a Goldilocks Girl, does not rule out

---

<sup>15</sup> For an in-depth discussion of this point, as it qualifies as evidence in favor of the suppositional view of subjunctive conditionals, see my forthcoming.

either of the two possibilities. This fits with a natural response to the question of whether such a girl would like candy: “There is no correct answer to the question. She might like candy; she might not.”

The fourth clue, then, is that competent users of ‘zif’ will attest to the fact that (8) does not apply to zif-statements. For illustration, let us suppose that there is no answer to the question whether, zif there were a Goldilocks Girl, she would like candy. In this context, we remain confident that such a girl *might* like candy and that such a girl *might not* like candy. This is further evidence that zif-statements are suppositional, rather than categorical, statements.

*Clue #5.* The fifth clue centers on the following principle:

(9a) However confident one is that *S*, one should be equally confident that it is true that *S*.

(9a) holds for all categorical statements, even those that involve a hidden subjective relativity. Suppose for instance that to state that *x* is *funny* is to state that *x* is funny *relative to the standards of the speaker*. However confident you are that Borat is funny (relative to your standards), you should be equally confident that the statement that Borat is funny (relative to your standards) is true. Generalizing, (9a) holds for all categorical statements.

(9a) does not hold for suppositional statements. Relative to a supposition *that A were the case*, it can be rational to be (i) highly confident *that C would be the case*, yet (ii) certain that it is neither true nor false *that C would be the case*. Relative to the supposition *that A were the case*, (i) commits one to the view that the supposition makes it highly probable that *C* would be the case, and (ii) commits one to the view that the supposition entails neither that *C* would be the case nor that *C* would not be the case. To repeat, there is nothing problematic about a supposition which makes something highly probable without entailing it or its negation, and (i) and (ii) have no further jointly problematic commitments. So (9a) does not hold for suppositional statements.

Competent users of ‘zif’ will attest to the fact that (9a) does not govern the rational evaluation of zif-statements. To be sure, we will endorse a similar principle:

- (9b) However confident one is that, *zif*  $A$  were the case,  $C$  would be the case, one should be at least as confident that, *zif*  $A$  were the case, a statement *that C was the case* would be true.

The fifth clue, then, is that competent users of ‘*zif*’ will attest to the fact that the rational evaluation of *zif*-statements is governed by (9b) rather than (9a). This clue reveals, once more, that *zif*-statements are suppositional, rather than categorical, statements.

Before considering the next clue, I want to pause to confirm a claim that I made in the preceding section. I said that hypothetical-scenario/possible-situation talk was equivalent to suppositional talk, and that because suppositional talk cannot be interpreted categorically, neither can hypothetical-scenario/possible-situation talk. To confirm that hypothetical-scenario/possible-situation talk cannot be interpreted categorically, notice that principles (7a), (8), and (9a) do not apply to hypothetical-scenario/possible-situation statements. For illustration, suppose that there were a Goldilocks Girl. In this scenario/situation, would such a girl like candy? One might naturally respond as follows: “Not necessarily. Of course, it is possible. Indeed, it is even highly probable, for most little girls like candy. It would be silly, however, to think that there is a fact of the matter. Strictly speaking, it is neither true nor false that, in this scenario/situation, such a girl would like candy. For although the description of the scenario/situation makes it likely that such a girl would like candy, it does not settle whether she would like candy.” The fact that (7a), (8), and (9a) do not apply to hypothetical-scenario/possible-situation statements reaffirms our earlier conclusion that these statements cannot be interpreted categorically.

*Clue #6.* Intuitively, it cannot be objectively incorrect to assign probability 1 to a categorical statement *and* objectively incorrect to assign probability 0 to the statement. If a categorical statement states a tautology, then it is not objectively incorrect to assign to it probability 1. If it states a contradiction, then it is not objectively incorrect to assign to it probability 0. And if it states neither a tautology nor a contradiction, then, while assignments of 1 and 0 may both be *subjectively* incorrect—that is, incorrect relative to a given epistemic situation—neither is *objectively* incorrect—that is, incorrect irrespective of one’s epistemic situation. Intuitively, then, assignments of probability 1 and 0 to a

categorical statement cannot both be objectively incorrect. (Standard interpretations of probability respect this intuition.)

By contrast, intuitively, assignments of probability 1 and 0 to a subjunctive suppositional statement *can* both be objectively incorrect. For to assign probability to a suppositional statement is to assign probability to what is stated by the statement, relative to what is supposed. And, intuitively, to assign probability 1/0 to what is stated by a subjunctive statement, relative to a subjunctive supposition, is to commit to the claim that what is stated is *entailed by/inconsistent with* what is supposed. Thus, for those subjunctive suppositional statements whose suppositions neither entail, nor are inconsistent with, what is stated by the statements, assignments of probability 1 and 0 will, intuitively, both be objectively incorrect.

For illustration, suppose that there were a Goldilocks Girl. Intuitively, to say that the probability is 1 that such a girl would like candy is to rule out the possibility that such a girl would not like candy, which is to commit to the claim that being a Goldilocks Girl entails liking candy. And to say that the probability is 0 that such a girl would like candy is to rule out the possibility that such a girl would like candy, which is to commit to the claim that being a Goldilocks Girl is inconsistent with liking candy. Because being a Goldilocks Girl seems neither to entail, nor to be inconsistent with, liking candy, it seems objectively incorrect to assign probability 1 or 0 to the statement that such a girl would like candy. More generally, it seems to be objectively incorrect to assign probability 1 or 0 to any subjunctive suppositional statement of something that is neither entailed by nor inconsistent with what is supposed. (Standard interpretations of probability do not respect this intuition.<sup>16</sup>)

---

<sup>16</sup> There are at least two reasons standard interpretations cannot respect the intuition.

First, standard interpretations are not even equipped to treat subjunctive suppositional statements. By now it should be clear that suppositional statements cannot be interpreted categorically. Thus, ascriptions of probability to such statements must be treated as ascriptions of conditional probability to what is stated by the statements, relative to what is supposed by them. But standard theories of probability define conditional probability in terms of absolute probabilities. And there is no such thing as the absolute probability of what is stated or of what is supposed by a *subjunctive* suppositional statement. For instance, there is no such thing as the absolute probability either of the supposition *that there were a Goldilocks Girl* or of the claim *that such a girl would like candy*. (For an attempt to address this difficulty within the

The sixth clue is that, to competent users of ‘zif’, if it seems that the antecedent of a zif-statement entails neither its consequent nor the negation of its consequent, it will seem objectively incorrect to assign probability 1 or 0 to the zif-statement. For instance, it seems objectively incorrect to say that it is 100% probable that, zif there were a Goldilocks Girl, she would like candy; and it seems objectively incorrect to say that it is 0% probable that, zif there were a Goldilocks Girl, she would like candy. This is further evidence that zif-statements are suppositional, rather than categorical, statements.

*Clue #7.* In the context of a subjunctive supposition, we commit to various relations between the supposition and what is stated by a subjunctive statement by evaluating the statement in various ways, for instance as *possible*, *probable*, or *definite*; we do not however commit to anything in particular by making a *bare* subjunctive statement. A bare suppositional statement is true on one condition (that what is stated be entailed by what is supposed) and false on a different condition (that what is stated be inconsistent with what is supposed). There is no single condition, *C*, such that the statement is true if *C* obtains and false if *C* does not obtain. Thus, a bare suppositional statement does not, by itself, signify that the world is in any given condition; it is in this sense insignificant. In practice, we signify that the world is in a given condition by qualifying suppositional statements—usually implicitly—for instance by *possibly*, *probably*, or *definitely*. A clue for the outsider is that, when the author of a bare zif-statement explicitly denies qualifying it, competent users of ‘zif’ will deem the statement insignificant in the preceding sense.

To begin to see what I have in mind, consider an analogy. Plausibly, statements of the form ‘The *F*s are *G*’ carry no specific commitment without implicit or explicit qualification by ‘a given number of’, ‘some of’, ‘most of’, ‘a significant number of’, or ‘all of’. In a context in which a speaker explicitly denies any such qualification, a competent audience will deem the statement insignificant. For instance:

---

framework of a standard interpretation of probability, see the first proposal for interpreting *Probability* in §2.)

Second, by the standard calculus, each of an infinite number of possibilities judged equiprobable has probability 0, and so an assignment of probability 0 to a proposition does not exclude the proposition as a possibility. But, intuitively, on the supposition that there were a Goldilocks Girl, an assignment of probability 0 to the claim that such a girl would like candy *does* exclude the possibility that such a girl would like candy.

Lucy: The cookies are stale.

Pedro: Do you mean that *all* of the cookies are stale?

Lucy: No. I don't mean that, nor do I believe it.

Pedro: Do you mean that *most*, or a *significant number*, of the cookies are stale?

Lucy: No. I do believe *that*, but it is not what I mean.

Pedro: Do you, then, mean that a *given number* of the cookies are stale?

Lucy: No.

Pedro: Do you merely mean that *some* of the cookies are stale?

Lucy: No. I believe that too, but it is not what I mean. When I say that the cookies are stale, I do not mean that *all*, *most*, a *significant number*, a *given number*, or *some* of the cookies are stale; I mean *that the cookies are stale*.

Pedro: Then I have lost all grasp of what it is that you are trying to tell me. By denying the various candidate qualifications, you have stripped your statement of any possible significance.

If an author of 'The Fs are G' explicitly denies any relevant qualification, a competent audience will deem the statement insignificant.

To see that a similar point applies to zif-statements, consider a variant of the original exchange between Ed and Ella:

Ed: Zif the glass had dropped, it would have shattered.

Ella: Do you mean that it *definitely* would have shattered?

Ed: No. I don't mean that, nor do I believe it.

Ella: Do you mean that it *probably* (or, very probably) would have shattered?

Ed: No. I do believe *that*, but it is not what I mean.

Ella: Do you merely mean that it *possibly* would have shattered?

Ed: No. I believe that too, but it is not what I mean. When I say that, zif the glass had dropped, it would have shattered, I do not mean that it *definitely*, *probably*, or *possibly* would have shattered; what I mean is that it *would* have shattered.

Ella: Then I have lost all grasp of what you are trying to tell me. By denying the various candidate qualifications, you have stripped your statement of any possible significance.<sup>17</sup>

The clue for the outsider centers on the fact competent users of ‘zif’ will deem explicitly bare zif-statements insignificant. But this is not the whole clue. For competent users of ‘the Fs’ deem explicitly bare statements of the form ‘The Fs are G’ insignificant and yet, when properly qualified, statements of this form are categorical. For the full clue, the outsider needs to consider the *sort* of qualification that zif-statements require. The candidate qualifiers can all be treated as sentential operators. Thus, what requires qualification is the thing that is stated by a zif-statement. From here the outsider can reason as follows: In order for a zif-statement to be significant, the thing that it states must be qualified. Things stated categorically do not require qualification for statements of them to be significant. For instance, there is no need to qualify the proposition *that Jones is the murderer*—by, say, ‘probably’, ‘definitely’, or ‘possibly’—in order for a categorical statement of it to be significant. By contrast, subjunctive contents stated relative to subjunctive suppositions *do* require qualification for their statements to be significant. Competent users of ‘zif’ attest to the fact that zif-statements require just the sorts of qualification that suppositional statements require. Hence, zif-statements are suppositional statements.

Before turning to the final clue, I want to consider a question that arises in light of the fact that explicitly bare zif-statements are, in the preceding sense, insignificant. The question concerns our practice of making bare zif-statements. Typically, one who makes a categorical statement aims to state something *true*. In other words, one aims to signify, of some condition that *obtains*, that the world is in that condition. But a bare zif-statement does not even signify that the world is in a given condition. So it seems that the aim of making a bare zif-statement cannot be to signify, of some condition that obtains, that the world is in that condition. What, then, is the typical aim of making a bare zif-statement?

---

<sup>17</sup> For the reader who doubts that Ed has stripped his statement of any possible significance, I suggest substituting ‘supposing’ for ‘zif’ in the conversation above and presenting it to non-philosophically-trained speakers of English. In my experience, all such speakers have agreed that the relevant statement has been stripped of significance. Indeed, the majority of them laugh when I say, “I do not mean that it *definitely, probably, or possibly* would have shattered; what I mean is that it *would* have shattered.”

In fact, it *is* to signify, of some condition that obtains, that the world is in that condition. The relevant difference between zif-statements and categorical statements lies not in the *aim* of making such statements, but rather in the *mode* of pursuing this aim. For categorical statements, the mode is direct: we simply state that the world is in a certain condition. For zif-statements, the mode is indirect: we qualify what is stated by the statement, relative to what is supposed, and we thereby incur a commitment to the world's being in a certain condition—a condition concerning the relationship between what is supposed and what is stated. With the utterance of a *bare* zif-statement, the qualification is *implicit*. Typically, the implicit qualification of choice is *probably*. When one qualifies what is stated by a zif-statement as *probable*, one commits to the idea that that what is stated is *made probable* by what is supposed. When we utter bare zif-statements, we typically aim to signify that world is in *this* condition. (Given my favored account of *Probability* in §2, we thereby aim to signify that certain connections and features of the world are *stable*.<sup>18</sup>)

It is worth noting that Gricean rules of conversation provide a common reason *not* to make explicit our qualification of zif-statements with ‘probably’. For illustration, consider another variant of the exchange between Ed and Ella:

Ed: Zif the glass had dropped, it would have shattered.

Ella: Not necessarily. A perfect gust of wind might have brought it to a gentle landing.

Ed: I did not say *necessarily*. Nor did I mean it. Obviously there is *some* chance that a perfect gust of wind might have brought it to a gentle landing. But that chance is insignificant. When I said that it would have shattered, I took it to be common knowledge that there was an *insignificant* chance that it would not have shattered. Given this knowledge, I thought it was potentially misleading for me to *explicitly* qualify my utterance with ‘probably’. For suppose that I had so qualified it. Then you could have been expected to reason as follows: “It is common knowledge that it is not certain that, zif the glass had dropped, it would have shattered. So there is no need for Ed to say that it *probably* would have shattered. All he needs to say is that it would have shattered. The rest is implicit. Why, then, did he explicitly qualify his statement by adding ‘probably’? He is a competent and cooperative speaker who does not make irrelevant contributions to the conversation. He must have added ‘probably’ for a reason. He must have intended to *emphasize* that it is merely probable, and to thus communicate that he has some reason to

---

<sup>18</sup> As Allan Gibbard puts it, in the case of subjunctive conditionals, we aim to “express our understanding of what depends on what in the world.” (1981, p. 239)

*doubt* that the glass would have shattered. I wonder what his reason is.” Because I have no reason to doubt that the glass would have shattered, it was potentially misleading for me to make explicit what was implicit all along: that I meant that it *probably* would have shattered.

Given that Gricean rules of conversation provide a common reason to leave implicit our qualification of zif-statements with ‘probably’, it is reasonable to expect bare zif-statements to be commonplace.

*Clue #8.* There are six equivalent ways of ascribing a (unary) value to a zif-statement, corresponding to the possible orderings of three components: the predicate used to ascribe the value, the that-clause used to denote the object of the evaluation (i.e., what is stated by the zif-statement), and the zif-clause used to express the supposition relative to which the value is assigned. For instance:

- (10a) It is probable that hamsters would fly, zif they had wings
- (10b) It is probable, zif hamsters had wings, that they would fly
- (10c) Zif hamsters had wings, it is probable that they would fly
- (10d) Zif hamsters had wings, that they would fly is probable
- (10e) That hamsters would fly, zif they had wings, is probable
- (10f) That hamsters would fly is probable, zif they had wings.

By contrast, ascriptions of values to many ordinary categorical statements admit of only two equivalent expressions, corresponding to the possible orderings of two components: the predicate used to ascribe the value, and the that-clause used to denote the object of the evaluation. For instance:

- (11a) It is probable that hamsters have feelings
- (11b) That hamsters have feelings is probable.

Competent users of ‘zif’ will attest to the fact that (10a) – (10f) are equivalent. The outsider may infer that ascriptions of values to zif-statements involve *three* main components: a predicate, a that-clause, and a zif-clause. The roles of the predicate and the that-clause are obvious: the predicate is used to ascribe

the value; the that-clause is used to denote the object of the evaluation. But what is the role of the zif-clause? The outsider needs to consider two hypotheses.

The first is that the zif-clause is used to *denote* something, perhaps a hypothetical situation, and the that-clause involves an implicit pronominalization on the zif-clause. This hypothesis is supported by the resemblance of zif/then-statements to both when/then- and where/there-statements. For illustration, consider the following statements (which I adapt from Lycan 2001):

(12a) When Sharon leaves, then I will leave. (When Sharon leaves, then, *at that time*, I will leave.)

(12b) Where Sharon lives, there I will live. (Where Sharon lives, there, *at that place*, I will live.)

(12c) Zif Sharon had left, then I would have left. (Zif Sharon had left, then, *in that situation*, I would have left.)

(12a) and (12b) appear to involve resumptive pronominalization on a denoting clause. In (12a), the clause ‘when Sharon leaves’ appears to denote a time; ‘then’ appears to resume this denotation. In (12b), the clause ‘where Sharon lives’ appears to denote a place; ‘there’ appears to resume this denotation. The surface resemblance of (12c) to both (12a) and (12b) suggests that the clause ‘zif Sharon had left’ denotes something, say, a hypothetical situation, and that ‘then’ resumes this denotation.<sup>19</sup> A further similarity among (12a) – (12c) is that ascriptions of values to any of them can be broken into three parts: a predicate, a that-clause, and a when-, where-, or zif-clause. For illustration, consider (13a) – (13d):

(13a) It is probable that I will live where Sharon lives

(13b) It is probable, where Sharon lives, that I will live

(13c) Where Sharon lives, it is probable that I will live

(13d) That I will live where Sharon lives is probable.

The remaining two combinations are awkward, to say the least:

(13e) Where Sharon lives, that I will live is probable

---

<sup>19</sup> For this interpretation of subjunctive conditionals, see Geis 1973 and Lycan 2001.

(13f) That I will live is probable, where Sharon lives.

Despite this awkwardness, the outsider might consider the hypothesis that, like when- and where- clauses, zif-clauses are denoting clauses, and that that-clauses in evaluations of zif-statements involve implicit pronominalization on the respective zif-clauses.

To see that the hypothesis is false, the outsider may consider whether, like when- and where- clauses, zif-clauses admit of apposition with descriptive phrases. The first part of the eighth clue is that competent users of ‘zif’ will attest to the fact that, whereas ‘the time when Sharon leaves’ and ‘the place where Sharon lives’ are grammatical, ‘the hypothetical situation zif Sharon had left’ is not. What best explains this difference is that, unlike when- and where-clauses, zif-clauses are not denoting phrases (recall that zif-clauses were designed specifically to avoid reference/denotation). The outsider who discovers this fact must reject the first hypothesis.

The second hypothesis is that the role of a zif-clause is to express—rather than denote—something, namely, a supposition. The outsider might reason as follows: To evaluate a suppositional statement is to evaluate what is stated by the statement, relative to what is supposed. Thus, there are three semantic components to a unary evaluation of a suppositional statement: (i) the object of evaluation; (ii) the value assigned to this object; and (iii) the supposition relative to which the assignment is made. To communicate one’s evaluation of a suppositional statement, one must communicate (i) – (iii). But one need not do so in any particular order. In their evaluations of zif-statements, speakers combine three components in no particular order. They use a that-clause to denote (i); a predicate to express (ii); and a zif-clause to express (iii). Thus, the role of a zif-clause, both in a zif-statement and in the evaluation of a zif-statement, is to express a supposition.

The eighth clue, then, is that there are six equivalent ways of ascribing a unary value to a zif-statement, corresponding to the possible orderings of three components: a predicate, a that-clause, and a zif-clause. Furthermore, the role of the zif-clause cannot be to denote something, for unlike when- and where-clauses, zif-clauses do not admit of apposition with descriptive phrases.

## 6. Zif would have been if

By stipulation, a sincere and literal utterance of ‘zif *A*’ is an expression of the supposition that *A*. In my view, the stipulated role of ‘zif’ is identical to the natural role of ‘if’. In other words, if our linguistic community had actually employed ‘zif’ according to our stipulation, ‘zif *A*’ would have meant what ‘if *A*’ actually means: *zif* would have been *if*.

My reasons for holding that *zif* would have been *if* are threefold. First, in every respect that I can think of, our practices with ‘zif’ match perfectly with our practices with ‘if’. Second, corresponding to the interpretive pitfalls of ‘zif’ are features of our practice with ‘if’ that, on the proposed suppositional interpretation, can be predicted to mislead theorists into holding just the sorts of categorical views that are commonly held about ‘if’—in particular, the various possible-worlds views. Third, for each of the last six interpretive clues that we identified in our practice with ‘zif’, there is a corresponding feature of our practice with ‘if’ that stands as evidence in favor of the suppositional interpretation of ‘if’.

Consider the following pairs of conditionals and suppositionals:

- (C1) If hamsters had wings, they would fly.
- (S1) Supposing that hamsters had wings, they would fly.
  
- (C2) If hamsters had wings, would they fly?
- (S2) Supposing that hamsters had wings, would they fly?

On the view that *zif* would have been *if*, conditionals (C1) and (C2) are simply more efficient ways of expressing what is expressed by the corresponding suppositionals; they are more efficient because they avoid reference to the contents of the suppositions they express. Standard views of conditionals are correct that (C1) and (C2) state and ask things; they are incorrect that (C1) and (C2) state and ask things *categorically*.

Some features of our practices with ‘if’ may give the impression that conditionals admit of categorical interpretations. For example, because the significance of a conditional statement depends both on what is stated and on what is supposed, in practice we never separate the contents of conditional

statements from their suppositional contexts. Indeed, in the case of subjunctives, it is impossible to give a complete report of what is stated without also reporting on what is supposed. When asked for instance what is stated by (C1), we respond by providing what is stated *together* with what is supposed: *that, if hamsters had wings, they would fly*. A theorist who is not actively considering a suppositional interpretation might think that we are providing only what is stated and might thus infer that (C1) states something categorically. Further pitfalls might lead the theorist specifically to the view that subjunctive conditionals are categorical statements of things about the nearest possible world or set of worlds.

Though pervasive, such interpretive pitfalls can be overcome. For our practices are rich with interpretive clues. Here are eight:

1. The explicit addition of a nearness-condition to a subjunctive conditional affects our evaluation of the conditional.
2. We judge certain subjunctive conditionals to be about impossible scenarios, and our confidence in such statements is sometimes low and sometimes high.
3. The rational evaluation of subjunctive conditionals is governed by (7c) rather than by (7a), which governs the rational evaluation of categorical statements:
  - (7a) However confident one is that *S*, one should be at least as confident that there is an answer to the question of whether *S*.
  - (7c) However confident one is that, if *A* were the case, *C* would be the case, one should be at least as confident that, if *A* were the case, there would be a question as to *whether C was the case*, which *would* have an answer.
4. The rational evaluation of subjunctive conditionals is not governed by (8), which governs the rational evaluation of categorical statements:
  - (8) On the supposition that there is no answer to the question whether *a* is *F*, one should have zero confidence that *a* might be *F*.
5. The rational evaluation of subjunctive conditionals is governed by (9c) rather than by (9a), which governs the rational evaluation of categorical statements:
  - (9a) However confident one is that *S*, one should be equally confident that the statement that *S* is true.
  - (9c) However confident one is that, if *A* were the case, *C* would be the case, one should be at least as confident that, if *A* were the case, a statement *that C was the case* would be true.

6. Intuitively, assignments of probability 1 and 0 to a categorical statement cannot both be objectively incorrect, yet assignments of probability 1 and 0 to a subjunctive conditional statement can both be objectively incorrect (whenever the antecedent entails neither the consequent nor its negation).
7. In order for a subjunctive conditional statement to be significant, the thing that it states must be qualified—by, say, ‘probably’, ‘definitely’, or ‘possibly’.
8. There are six equivalent ways to ascribe a unary value to a subjunctive conditional statement, corresponding to the possible orderings of three components: the predicate used to ascribe the value, the that-clause used to denote the object of the evaluation, and the if-clause. Furthermore, the role of the if-clause cannot be to denote something, for unlike when- and where-clauses, if-clauses do not admit of apposition with descriptive phrases.

I have omitted two important topics from my treatment of conditionals. First, how should we treat compound statements with conditional statements as parts? Second, how should we characterize valid reasoning with conditionals? For indicative conditionals, I treat these questions in some detail in my 2006. (For similar treatments, see Adams 1965, 1975, 1998, and Edgington 1995; for discussion, see Bennett 2003.) Because my strategy for treating subjunctives is generally the same, for the sake of brevity, I refer the reader to my 2006.

Anyone who rejects that *zif* would have been *if* faces the obvious challenge: to find a relevant difference between our entrenched practices with ‘if’ and our inchoate practices with ‘zif’.<sup>20</sup>

## REFERENCES

- Adams, E. W. 1965. “A Logic of Conditionals”. *Inquiry*, 8, pp. 166–97.
- 1966. “Probability and the Logic of Conditionals”, in J. Hintikka and P. Suppes (eds.) 1966, pp. 256–316.
- 1975. *The Logic of Conditionals*. Dordrecht: Reidel.

---

<sup>20</sup> For helpful comments I am grateful to Christian Lee, Brad Monton, Adam Pautz, and an anonymous referee for *Noûs*.

- 1976. “Prior Probabilities and Subjunctive Conditionals”, in W. L. Harper and C. A. Hooker (eds.) *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, Vol. I, D. Reidel, Dordrecht, Holland, 1976, pp. 1-21.
- 1998. *A Primer of Probability Logic*. Stanford, CA: CSLI Publications.
- Author. 2006.
- . Forthcoming.
- Barker, S. 1999. Counterfactuals, Probabilistic Counterfactuals, and Causation. *Mind* 108: 427-469.
- Bennett, J. 1984. “Subjunctives and Temporal Direction”. *Philosophical Review*, 93, pp. 57-91.
- 2003. *A Philosophical Guide to Conditionals*. New York: Oxford.
- Black, M. (ed.) 1950. *Philosophical Analysis*. Englewood Cliffs: Prentice-Hall.
- Brown, R. and Rollins, C. D. (eds.) 1969. *Contemporary Philosophy in Australia*. London: Routledge.
- Chellas, B. 1975. “Basic Conditional Logic”. *Journal of Philosophical Logic*, 4, pp. 133-228.
- Cole, P. (ed.) 1978. *Syntax and Semantics, Volume 9: Pragmatics*. New York: Academic Press.
- Dudman, V. H. 1994. “On Conditionals”. *Journal of Philosophy*, 91, pp. 113–28.
- Dummett, M. 1978. *Truth and Other Enigmas*. London: Duckworth.
- Edgington, D. 1986. “Do Conditionals Have Truth Conditions?” in Jackson, F. (ed.) 1991, pp. 176–201.  
First published in *Critica* 18, 52, 1986, pp. 3–30.
- 1995. “On Conditionals”. *Mind*, 104, pp. 235–330.
- Eells, E. and Skyrms, B. (eds.) 1994. *Probability and Conditionals*. Cambridge: Cambridge University Press.
- Ellis, B. 1969. “An Epistemological Concept of Truth” in Brown, R. and Rollins, C. D. (eds.) 1969, pp. 72–88.
- Farley, A. M., Farley, P., McCullough, K. E. (eds.) 1986. *Papers from the Parasession on Pragmatics and Grammatical Theory, Twenty-Second Regional Meeting of the Chicago Linguistic Society*. Chicago: University of Chicago.
- Field, H. 2000. “Indeterminacy, degree of belief, and excluded middle”. *Nous* 34: 1-30.

- French, P., Uehling, T., and Wettstein, H. (eds.) 1977. *Midwest Studies in Philosophy* Volume 2, Minneapolis, MN: University of Minnesota Press.
- Gärdenfors, P. 1986. "Belief Revision and the Ramsey Test for Conditionals". *Philosophical Review*, 95, pp. 81–93.
- 1988. *Knowledge in Flux*. Cambridge, MA: MIT Press.
- Geis, M.L. 1973. "If and Unless". In Kahane, H., Kachru, B., and Kahane, R. (eds.) 1973, pp. 21–37.
- Gibbard, A. 1981. "Two Recent Theories of Conditionals". In Harper, W. L., Stalnaker, R., and Pearce, C. T. (eds.) 1981, pp. 211–47.
- Goodman, N. 1947. "The Problem of Subjunctive Conditionals". *Journal of Philosophy* 44: 113-28.
- Grice, H.P. 1967. "Indicative Conditionals". In Grice 1989, pp. 58–85.
- 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Harper, W. L., Stalnaker, R., and Pearce, C. T. (eds.) 1981. *Ifs*. Dordrecht: Reidel.
- Hazen, A. and Slote, M. 1979. "Even if". *Analysis*, 39, pp. 35–41.
- Hiddleston, E. 2005. A Causal Theory of Counterfactuals. *Noûs* 39: 632-657.
- Hintikka, J. and Suppes, P. (eds.) 1966. *Aspects of Inductive Logic*. Amsterdam: North Holland.
- Jackson, F. 1979. "On Assertion and Indicative Conditionals". *Philosophical Review*, 88, pp. 565–89.
- 1987. *Conditionals*. Oxford: Blackwell.
- (ed.) 1991. *Conditionals*. Oxford: Clarendon Press.
- Jeffrey, R. 1991. "Matter of Fact Conditionals". *Proceedings of the Aristotelian Society Supplementary*, 65, pp. 161–83.
- Jeffrey, R. and Stalnaker, R. 1994. "Conditionals as Random Variables" in Eells, E. and Skyrms, B. (eds.) 1994, pp. 49–66.
- Kahane, H., Kachru, B., and Kahane, R. (eds.) 1973. *Issues in Linguistics: Papers in Honor of Henry and Renee Kahane*. Urbana, Ill.: University of Illinois Press.
- Kaplan, D. 1978. "Dthat" in Cole, P. (ed.) 1978, pp. 221–43.
- Kment, B. 2006. "Counterfactuals and the Analysis of Necessity". *Philosophical Perspectives*, 20, pp. 237-302.

- Kratzer, A. 1986. "Conditionals" in Farley, A. M., Farley, P., McCullough, K. E. (eds.) 1986, pp. 1–15.
- Kripke, S. 1977. "Speaker's Reference and Semantic Reference" in French, P., Uehling, T., and Wettstein, H. (eds.) 1977, pp. 6–27.
- Lewis, D. 1973. *Counterfactuals*. Cambridge, Mass: Harvard University Press.
- 1976. "Probabilities of Conditionals and Conditional Probabilities". *Philosophical Review*, 85, pp. 297–315. Page references to Lewis 1986.
- 1986. *Philosophical Papers* Volume 2. Oxford: Oxford University Press.
- 1986(a). "Postscripts to Probability of Conditionals and Conditional Probabilities" in Lewis, D. 1986, pp. 152–56.
- Lycan, W. 2001. *Real Conditionals*. New York: Oxford University Press.
- 1991. "Even and Even if". *Linguistics and Philosophy*, 14, pp. 115-50.
- Mackie, J. 1973. *Truth, Probability and Paradox*. Oxford: Clarendon Press.
- McDermott, M. 1996. "On the Truth Conditions of Certain 'If'-Sentences". *Philosophical Review*, 105, pp. 1–37.
- McGee, V. 1985. "A Counterexample to Modus Ponens". *Journal of Philosophy*, 82, pp. 462–71.
- 1989. "Conditional Probabilities and Compounds of Conditionals". *Philosophical Review*, 98, pp. 485–542.
- Nolan, D. 1997, "Impossible Worlds: A Modest Approach". *Notre Dame Journal for Formal Logic*, 38, pp. 535-572.
- Pollock, J. L. 1976. *Subjunctive Reasoning*. Boston: Reidel.
- Quine, W. V. O. 1952. *Methods of Logic*. London: Rutledge and Kegan Paul. Page references to third edition, 1974.
- Ryle, G. 1950. "If, 'So' and 'Because'", in Black, M. (ed.) 1950.
- Skyrms, B. 1981. "The Prior Propensity Account of Subjunctive Conditionals". In W. Harper, G. Pearce, and R. Stalnaker, 1981, pp. 259–65.
- 1994. "Adams Conditionals". In Eells, E. and Skyrms, B. (eds.) 1994, pp. 13-26.
- Slote, M. A. 1978. "Time in Subjunctives". *Philosophical Review*, 87, pp. 3-27.

- Stalnaker, R. 1968. "A Theory of Conditionals". In *Studies in Logical Theory, American Philosophical Quarterly Monograph Volume 2*, pp. 98–112.
- 1970. "Probability and Conditionals". *Philosophy of Science*, 37, pp. 64–80.
- 1975. "Indicative Conditionals". *Philosophia*, 5, pp. 269–86.
- . 1981. "A Defense of Conditional Excluded Middle". In W. Harper, G. Pearce, and R. Stalnaker, 1981, pp. 87–104.
- Stalnaker, R. and Thomason, R. 1970. "A Semantic Analysis of Conditional Logic". *Theoria*, 36, pp. 23–42.
- Williamson, T. 1996. "Knowing and Asserting". *Philosophical Review*, 105, pp. 489-523.
- Woods, M. 1997. *Conditionals*. Oxford: Clarendon Press.