Subjunctivity and Conditionals

Kai F. Wehmeier*

I take it that a satisfactory logical analysis of English *if* sentences should fulfill the following three desiderata.

(A) It should predict correct truth conditions for as wide a class of English conditionals as possible.

(B) It should illuminate the interaction of *if* with the indicative and subjunctive[1] moods in all grammatically admissible configurations.

(C) It should explain how the single English word *if* is able to perform the ostensibly distinct tasks it does in indicative and subjunctive conditionals.

* This paper is part of a larger research project on the logical role of the indicative-subjunctive distinction that began with the introduction of the subjunctive modal logic SML in (Wehmeier 2004). The subjunctive-indicative conditionals adduced in §2 bear obvious similarities to the examples I use to motivate SML there and in (Wehmeier 2005). In the latter work (p. 196) I use the subjunctive-indicative conditional 'If everyone who is rich had been poor, then someone who is poor would have been rich' as an example and suggest that 'it is even clearer in the case of counterfactual conditionals than in the case of the modalities that it is the interplay between the conditional operator and the subjunctive that effects extension-shifting to other possible worlds.' Other projects, including an investigation of the role of subjunctivity in cross-world predication (Wehmeier forthcoming), temporarily kept me from pursuing the issue of subjunctive conditionals. The present essay grew out of a much shorter paper on conditionals that I circulated in June 2010. John Burgess, Kai von Fintel, Lloyd Humberstone, and Helge Rückert provided very helpful comments on that paper, leading in particular to a much more satisfactory treatment of indicative conditionals. Moreover, I am indebted to Burgess for suggesting example (5) as a test case for my account, to von Fintel for pressing me on the question how indicative predicates can invariably pertain to the actual world, to Humberstone for pointing me to the insightful paper (Weatherson 2001), and to Rückert for suggesting a way of handling nested subjunctive conditionals (that has not found its way into the present essay).

[1] I use 'subjunctive' as a generic term covering both the 'fake past' (Iatridou 2000) of the antecedents of subjunctive conditionals and the 'conditional mood' of their consequents. A better choice might have been 'irrealis', but it seems advisable to follow the entrenched terminology. In any case, as Burgess (2011: 7-8) points out, the use of different grammatical moods for the antecedents and consequents, respectively, of subjunctive conditionals seems to be "an historically comparatively recent development in English". Burgess cites the example "... if thou hadst been here, my brother had not died" (John 11:22 in the King James version) as evidence that the subjunctive used to be used in both antecedent and consequent position. German still uses the subjunctive mood (*Konjunktiv II*) in both places, as did Latin (*coniunctivus irrealis*).

In this essay, I show that neither Robert Stalnaker's (1968; 1965) nor David Lewis's (1973) possible-worlds accounts of conditionals is fully satisfactory with respect to these desiderata, and I outline a new account that does substantially better.

The plan is as follows:

In §1 I offer a brief sketch of Stalnaker's and Lewis's accounts.

I then, in §2, introduce a type of conditional, to be called subjunctive-indicative, that has not received much attention in the literature[2] but plays a central role in the argument of this paper. The section continues by showing that Stalnaker's and Lewis's accounts predict incorrect truth conditions for such subjunctive-indicative conditionals, which establishes that they fall short with respect to desideratum (A). Given that both accounts offer explanations of the rôle of the indicative and subjunctive moods that, in light of subjunctive-indicative conditionals, cannot be correct, it follows that they also fail with respect to desideratum (B). Lewis (1973, 1976) assumes that there are semantically distinct indicative and subjunctive conditional operators (thus, in effect, postulating lexical ambiguity of *if*), and so his theory does not even attempt to abide by desideratum (C). Stalnaker, by contrast, provides a *prima facie* attractive explanation based on pragmatic differences between indicative and subjunctive conditionals. But his pragmatic distinction turns out to be insufficient to account for the truth conditions of subjunctive-

[2] An exception of sorts is (Percus 2000), who discusses conditionals similar to example (3) below, though not with respect to the indicative-subjunctive distinction.

indicative conditionals, and so Stalnaker's theory, too, ultimately falls short with respect to desideratum (C).

In §3, I present my own account of subjunctive (and subjunctive-indicative) conditionals, which is based on a broadly Lewisian possible-worlds framework (*sans* counterparts) but locates indicativity and subjunctivity in individual predicates rather than at the level of the conditional construction. Roughly, indicative predicates are invariably interpreted at the actual world, whereas subjunctive predicates pertain to the world of evaluation. I conclude the section by showing that the proposed account takes subjunctive-indicative conditionals in stride and coincides with Lewis's analysis for the special case of purely subjunctive conditionals.

The following §4 is devoted to an integration of indicative conditionals into the framework of §3. A potential problem originating from the treatment of indicativity—the threat of having indicative conditionals collapse into material ones—is averted by making use of an idea of Brian Weatherson's (2001), according to which indicative conditionals get their epistemic feel because their evaluation requires variation of which world is considered as actual, rather than of the world of evaluation.[3] While I am not wedded to the details of this particular treatment of indicative conditionals, I take it to be crucial that it establishes how (a) a single operator can play the epistemic role required for indicative conditionals and the metaphysical role required for subjunctive ones, and (b) my

[3] The idea can ultimately be traced back to Martin Davies and Lloyd Humberstone (1980) and Gareth Evans (1979), although it is used here in a framework much more similar to that of (Stein 2011). It is the origin of the research program known as two-dimensional semantics. See (Humberstone 2004).

insistence that indicative predicates always be evaluated at the world considered as actual does not trivialize indicative conditionals.

In §5 I consider ways of amending Stalnaker-Lewis type theories by adding an actuality operator to the conditional language. I argue that this results either (a) in a system that predicts an unobservable ambiguity in the modal profiles of ordinary indicative utterances, or (b) in a system whose treatment of English indicative and subjunctive predicates is ad hoc, or (c) in a notational variant of the language proposed in §3.

The final §6 provides a synopsis of the account developed in this paper and explains how it satisfies desiderata (A) through (C).

1. The Received View

Both Stalnaker and Lewis assume that conditionals fall into two broad categories, subjunctive and indicative. I begin by outlining their views about the former. Consider by way of example:

(1) If Shakespeare hadn't written *Hamlet*, somebody else would have.

Stalnaker and Lewis both assign (1) the logical form

(1a) $A > C$,

with A standing for ‘Shakespeare didn’t write *Hamlet*’ and C for ‘Somebody other than Shakespeare wrote *Hamlet*.’ The truth conditions they postulate for sentences of the form (1a) are quite similar, if not identical. Glossing over such niceties as Lewis’s rejection of the limit assumption, we can say that they both take $A > C$ to be true just in case all worlds closest to the actual world in which A is true are worlds in which C is true (for short: the closest A-worlds are C-worlds). The accounts differ in that Lewis allows for the possibility that there are several closest A-worlds, while Stalnaker insists that there always be exactly one (the ‘uniqueness assumption’).

The subjunctive conditional (1) contrasts with its indicative cousin

(2) If Shakespeare didn’t write *Hamlet*, somebody else did.

Stalnaker and Lewis disagree over the proper account of the difference between (1) and (2). Lewis takes (2) to be of the form

(2a) $A \rightarrow B$,

where ‘$\rightarrow$’ stands for a conditional operator semantically distinct from ‘$>$’. According to Lewis, the arrow ‘$\rightarrow$’ is to be read simply as ‘if’, while the corner ‘$>$’ is to be glossed as ‘if it were the case that... it would be the case that...’. Thus, Lewis takes the instantiation of (1a) to (1) to proceed in two stages. First one forms ‘If it were the case that

Shakespeare didn't write *Hamlet*, it would be the case that somebody else did,' and then one amalgamates the subjunctive mood contained in the operator '>' with the indicative mood of the predicates in the antecedent and consequent in order to obtain (1) itself.[4] But Lewis has little to say about the arrow; he seems to have been content to regard it as standing for the material conditional. It does not matter for our purposes what Lewis thought the correct semantics for '→' was; we simply note that, on his view, the English connective *if* is lexically ambiguous between the arrow and the corner.

While Lewis's explanation of the difference between (1) and (2) is semantic, Stalnaker's is pragmatic. According to Stalnaker, (1a) is the logical form of both (1) and (2), but the indicative mood in (2) signals the presence of a certain pragmatic constraint, viz. that the closest antecedent-world must be picked from among those compatible with the presuppositions that are in place at the time (2) is uttered.[5] The subjunctive mood of (1), by contrast, indicates the absence of any such constraint. Stalnaker can thus explain the difference in truth value between (1) and (2) by noting that nobody wrote *Hamlet* in the closest possible world in which Shakespeare didn't, so that (1) is false; but in the closest possible world *in the context set* in which Shakespeare didn't write *Hamlet*, somebody else did write the work— because its existence is a presupposition underlying the utterance of (2). Hence (2) is true.

This difference between Stalnaker and Lewis notwithstanding, their accounts agree that indicativity and subjunctivity are features that apply at the level of the conditional

---

[4] Cf. (Lewis 1973, 1-3) for an explicit description of this procedure, as well as (Bennett 2003, 5-6) and (Nute and Cross 2002, 2) for similar views.

[5] Worlds that are so compatible are said to belong to the context set.

construction, not at the sub-sentential level. To isolate, for instance, the predicate ‘would have written *Hamlet*’ and ask after its contribution to the truth-conditions of (1) is misguided on either view. According to Lewis, ‘would’ derives from the subjunctive conditional operator but does not attach to the predicate in any semantically relevant sense. According to Stalnaker, ‘would’ is part of what makes the entire conditional subjunctive, thereby signaling the absence of the context set requirement; but this has nothing to do with the predicate *per se*, only with the mechanism for evaluating the entire conditional.

Let us call an English conditional purely subjunctive (indicative) if all predicates in its antecedent and consequent are subjunctive (indicative). Stalnaker’s and Lewis’s accounts are geared toward conditionals that are either purely subjunctive, such as (1), or purely indicative, like (2). There are, however, perfectly innocuous English *if* sentences that do not fit either mold, as we will see in the following section.

2. Subjunctive-Indicative Conditionals

Let us assume that the Department of Philosophy at the University of Afla Ateb has eleven members, among them Anna, Barbara, and Christa, all three of whom are running for department chair. Each faculty member has exactly one vote, and a plurality suffices for electoral victory. When the votes are tallied, it turns out that Anna garnered five, and Barbara and Christa three votes each. David, not a big fan of Anna’s leadership, grumbles:

(3) If everybody who voted for Christa had voted for Barbara, Anna wouldn't have been elected chair.

This is true, for if the three faculty who opted for Christa had instead lent their support to Barbara, Barbara would have obtained six votes, thus prevailing over Anna with her five.

What is remarkable about (3) is the presence of both a subjunctive ('had voted') and an indicative ('voted') predicate in the antecedent. To be more precise, the main verb of both apodosis and protasis is subjunctive, but the protasis contains an indicative verb as well.

To continue with the goings-on at the University of Afla Ateb, suppose Barbara promised that, if elected, she would get rid of the clutter in the library by throwing out some of the logic books. Her plans weren't specific—she was going to pick a sufficient number of logic texts at random. A few months after the election, our friend David sees an eager student check out all the logic books. Then David (who seems to have a penchant for subjunctive talk) speaks truly when he says:

(4) If Barbara had been elected chair, some of the books you are checking out would have been thrown away by now.

Again, the main verbs of antecedent and consequent are in subjunctive mood. This time, however, it is the consequent that also contains an indicative verb.

It will prove helpful to introduce some terminology. Call an English *if* sentence subjunctive (indicative) if the main verbs of its protasis and apodosis are subjunctive (indicative). A subjunctive conditional is subjunctive-indicative if its protasis or its apodosis (or both) contains an indicative predicate. Thus (3) and (4) are subjunctive-indicative conditionals, while (1) is purely subjunctive; all three conditionals are subjunctive *tout court*.

Let us consider how Stalnaker's and Lewis's theories, respectively, might handle the subjunctive-indicative conditional (4).

Take Lewis's account first. The fact that (4) contains an indicative predicate (to wit, 'are checking out') in its apodosis implies, according to Lewis, that the grammatical consequent of (4) cannot be its logical consequent. For recall how Lewis's formalization procedure works: Indicative phrases are plugged into the antecedent and consequent positions of the operator 'if it were the case that... it would be the case that...', and then the subjunctive mood of the conditional construction is redistributed into the verbs of the antecedent and consequent. This makes it impossible for indicative predicates to occur within the scope of a Lewisian subjunctive conditional. We are forced to conclude that the indicative predicate 'are checking out' falls *outside* the logical consequent of the subjunctive conditional operator. In other words, Lewis's account yields the following analysis:

(4a) $\exists y\,[\mathrm{C}y \,\&\, (\mathrm{E}b > \mathrm{T}y)]$,

where the bound variable *y* ranges over the volumes in the philosophy library, *b* stands for 'Barbara,' C*y* for 'you are checking out *y*', E*x* for '*x* was elected chair,' and T*y* for '*y* has been thrown away'. Drawing the indicative 'are checking out' out of the consequent of the conditional in this manner saves it from being subjunctivized by the corner operator. Retranslated into English, (4a) reads: 'At least one of the books you are checking out is such that, had Barbara been elected chair, it would have been thrown away by now.'

But on Lewis's account, (4a) has the wrong truth-value. Recall that there is no *particular* volume that Barbara intended to throw out: For any given logic book, there is an E*b*-world in which that book survives Barbara's purge. Thus we cannot point to any one book that, had Barbara become chair, would have been chucked; in other words, there is no witness to the leading existential quantifier in (4a). Hence (4a), unlike (4), is false.

Before turning to Stalnaker, let us ask what went wrong for Lewis. It seems that (4a), while derived from Lewis's account of indicative and subjunctive mood, erroneously represents (4) as an existentially quantified statement, when in fact its main operator is the conditional: The existential witness should depend on the possible world in which the antecedent holds; it need not be specified independently. This suggests that we might be able to salvage Lewis's account by representing (4) as

(4b) $Eb > \exists y\,(Cy \,\&\, Ty)$,

in words: ‘Had Barbara become chair, some of the books that you would be checking out would have been thrown away’. But this will not work. After all, it’s impossible to check out a book that’s been thrown away.[6]

What are Stalnaker’s options for treating (4)? It should be clear that (4b) delivers the wrong truth conditions, no matter whether we apply Lewis’s or Stalnaker’s theory. This leaves us with

(4a) $\exists y\,[Cy \,\&\, (Eb > Ty)]$,

which we rejected when considering Lewis’s theory. Somewhat surprisingly, perhaps, this formalization does have the right truth-value on Stalnaker’s account—if only because of the (contentious) uniqueness assumption. For let w be the (unique) closest Eb-world, and let l be one of the logic books that Barbara throws out in this particular world w. Then l can serve as an existential witness for (4a). To be sure, the uniqueness assumption appears completely implausible in our scenario, and intuitively, one would want to regard the natural language equivalent of (4a), ‘At least one of the books you are checking out is such that, had Barbara been elected chair, it would have been thrown away,’ as false. It is thus not at all clear that Stalnaker’s account really has a leg up on Lewis’s with respect to (4), but at least he can provide a formalization that has the right truth-value.

[6] Note also that the retranslation of (4b) into English gets the moods wrong.

Let's consider our other example:

(3) If everybody who voted for Christa had voted for Barbara, Anna wouldn't have been elected chair.

Again, on Lewis's account, the indicative predicate 'voted for Christa' cannot lie within the logical antecedent of (3), and thus this sentence must be formalized as

(3a) $\forall x\ [\mathrm{V}xc \supset (\mathrm{V}xb > {\sim}\mathrm{E}a)]$,

where the bound variable 'x' ranges over the voting members of the department, '$\mathrm{V}xy$' stands for 'x voted for y', and 'a' and 'c' for 'Anna' and 'Christa', respectively.[7] Retranslated into English, (3a) reads: 'Everybody who voted for Christa is such that, had he voted for Barbara, Anna wouldn't have been elected chair.'

However, (3a) cannot be an adequate formalization of (3), whether we are using Lewis's or Stalnaker's theory. For let us assume David voted for Christa. Then it follows from (3a) that, if David had voted for Barbara, Anna would not have been elected chair. Not so: If David had voted for Barbara, Barbara would have garnered four votes instead of three, but Anna would still have had five votes, and would therefore have won the

[7] I use the horseshoe '$\supset$' for the material conditional. Readers who find the presence of the material conditional in (3a) distracting may wish to consider instead the formalization $[\forall x\text{: } \mathrm{V}xc]\ (\mathrm{V}xb > {\sim}\mathrm{E}a)$, which uses a binary quantifier.

election.[8] The problem with the Lewisian analysis seems to be, again, that it makes out (3) as a quantified statement, when in fact it has the logical form of a conditional.

Let us turn to Stalnaker's analysis. Here, too, one might have thought (3a) the formalization of choice, but we have already seen that it is inadequate. Stalnaker, however, could try to treat (3) as a conditional. In that case, its logical form must be

(3b) $[\forall x\ (\mathrm{V}xc \supset \mathrm{V}xb)] > {\sim}\mathrm{E}a$.[9]

A moment's thought, however, shows that (3b) also delivers the wrong truth-value. Given that every faculty member has only one vote, the antecedent $\forall x\ (\mathrm{V}xc \supset \mathrm{V}xb)$ of (3b) is equivalent, in every world reasonably close to the actual one, to Christa's obtaining no votes whatsoever, i.e. to $\forall x\ {\sim}\mathrm{V}xc$ (because $\mathrm{V}xc$ and $\mathrm{V}xb$ cannot jointly be true, given the voting rules). But then (3b) has the same truth conditions as

$(\forall x\ {\sim}\mathrm{V}xc) > {\sim}\mathrm{E}a$,

in words:

If nobody had voted for Christa, Anna wouldn't have been elected chair.

---

[8] The same consideration shows that the existential sentence $\exists x\ [\mathrm{V}xc\ \&\ (\mathrm{V}xb > {\sim}\mathrm{E}a)]$ also fails as a formalization of (2).

[9] This, too, can be formalized with the help of a binary quantifier, rather than a material conditional, thus: $([\forall x\!: \mathrm{V}xc]\ \mathrm{V}xb) > {\sim}\mathrm{E}a$.

This, unlike (3), is false: Suppose, for instance, that everybody who voted for Christa had voted for Anna. Then Anna would have garnered eight votes and thus scored a landslide victory.

It is easy to see what causes the trouble: On Stalnaker's theory, subjunctivity and indicativity can only be attributed to entire conditionals, not selectively to constituent predicates, and thus formalizing (3) as a conditional forces the two predicates in the antecedent to be evaluated in the same world, despite the fact that one of them is indicative (and should thus pertain to the actual world).

Time to take stock.

Lewis's account founders in the face of conditionals (3) and (4). It predicts incorrect truth-values for them, and it does so because of its particular assumption about the connection between the corner operator and the grammatical moods of the predicates occurring within its scope. Desiderata (A) and (B) are therefore not satisfied by Lewis's theory.

On one reading, Stalnaker's account yields the right truth-value for (4), although this seems to be due to a quirk of his analysis (the uniqueness assumption). Be that as it may, Stalnaker definitely predicts incorrect truth conditions for (3), and so his theory, too, fails to satisfy desideratum (A). As discussed, this is due to the theory's inability to apply the indicative-subjunctive distinction at a subsentential level: the indicative predicate 'voted' in (3) should be evaluated at the actual world, which is impossible if it occurs within the scope of a subjunctive Stalnaker conditional. This shows that Stalnaker's account falls short with respect to desideratum (B) as well.

What about (C), the requirement that an account of conditionals explain how English *if* can function in such apparently diverse ways as it does in indicative and subjunctive conditionals? Lewis's theory, as we have observed, simply gives up on this issue, postulating lexical ambiguity. Stalnaker's view is more subtle and at least purports to provide the desired explanation: There is only one English connective *if*; it is just that, when *if* is used with constituent sentences in indicative mood, the selection of an antecedent-satisfying world is pragmatically constrained by the context set requirement, whereas no such constraint is in place when protasis and apodosis are in subjunctive mood.

This is an attractive view, as far as it goes. However, making the distinction between indicative and subjunctive mood only at the level of the conditional operator lies at the root of the intractability of (3) in Stalnaker's account. For when we reach the corner in an evaluation of

(3b) $[\forall x\ (\mathrm{V}xc \supset \mathrm{V}xb)] > {\sim}\mathrm{E}a$,

we have already interpreted both occurrences of 'V' in the antecedent according to the *same* world of evaluation, when in fact the first occurrence ought to be interpreted in the actual world. Attending to subjunctivity only at the level of the conditional leads to picking the wrong antecedent-satisfying world, and consequently such an explanation for the uniform meaning of *if* is ultimately unsuccessful. Desideratum (C) thus proves a hurdle for Stalnaker's account as well.

3. Subjunctivity and Conditionals: A New Account

As noted, the assumption that indicativity and subjunctivity apply to a conditional as a whole is responsible for the failure of Stalnaker's and Lewis's accounts *vis-à-vis* subjunctive-indicative conditionals. I therefore propose to renounce this assumption, and to regard indicativity and subjunctivity as semantic properties of individual predicates.

I submit that subjunctive predicates are *bound* to the conditional operator in whose immediate scope they occur, in the sense that they are evaluated at the possible worlds introduced by that operator. Indicative predicates, by contrast, bear no such binding relation to the conditional operator, and consequently pertain to the actual world no matter where they occur.

More specifically, I suggest that (3) has the following logical form:

(3s) $[\forall x\ (\mathrm{V}xc \supset \mathrm{V}^{s}xb)] \Rightarrow {\sim}\mathrm{E}^{s}a$, where the superscript ‘s’ (the *subjunctive marker*) indicates subjunctive mood (and absence of subjunctive marker, indicative mood), and the double arrow ‘$\Rightarrow$’ represents the *if* of (3).[10, 11]

Similarly, I hold that (4) has the form

(4s) $\mathrm{E}^{s}b \Rightarrow \exists y\ (\mathrm{C}y\ \&\ \mathrm{T}^{s}y)$.

Presently I will turn to the task of outlining a semantics for a language with indicative and subjunctive predicates and the conditional operator $\Rightarrow$, but first we must address a point emphasized by Dudman (1984).

Dudman’s observation is that neither the antecedents nor the consequents of (3) and (4), taken *verbatim*, appear to express propositions, as they cannot on their own be used to make an assertion (or at least not an assertion with the requisite content). For example, ‘Anna wouldn’t have been elected chair’, standing by itself, fails to express a proposition because no counterfactual scenario for the evaluation of the subjunctive predicate ‘would have been elected chair’ has been specified.[12] The fact that the antecedents and consequents of our formal representations (3s) and (4s) are not formulas of first-order

---

[10] Those who prefer a binary quantifier should read (3s) as $([\forall x\colon \mathrm{V}xc]\ \mathrm{V}^{s}xb) \Rightarrow {\sim}\mathrm{E}^{s}a$.

[11] Humberstone (1982) first proposed using a syntactic device (in his case, a sentential operator) corresponding to subjunctive mood in modal logic. In this same context, Wehmeier (2004) first introduced the subjunctive marker as used here. Barker (1996) takes a step in this direction for subjunctive conditionals but stops halfway: he applies his mood operator to the entire antecedent of a conditional, a move that prevents selective subjunctivization of individual predicates.

[12] Similarly, ‘everybody who voted for Christa had voted for Barbara’, when considered by itself, at best says that all those who, at some unspecified time in the past, voted for Christa, at some earlier time voted for Barbara, but this is clearly not the meaning the antecedent of (2) contributes to the meaning of the conditional.

logic, but rather have the subjunctive marker ‘s’ attached to some of their predicates, underscores this observation.

The correctness of Dudman’s point notwithstanding, it does not support the further claim that (3) and (4) are not, properly speaking, composed of what (3s) and (4s) present as their antecedents and consequents. That ‘Anna wouldn’t have been elected chair’ cannot be used, by itself, to make an assertion is no more worrisome than that ‘she mocks him’ cannot so be used without salient referents for the pronouns. Nobody would take the latter fact to support an argument that ‘she mocks him’ is not a proper constituent of ‘Every man who has a daughter loves her even when she mocks him’, where the requisite referents are of course provided by the pronoun-binding quantifiers. In the same way, I suggest, both ‘everybody who voted for Christa had voted for Barbara’ and ‘Anna wouldn’t have been elected chair’, as occurring in (3), have the requisite counterfactual worlds provided by the subjunctive-binding conditional operator. Indeed, the subjunctive mood itself essentially *is* a bound world-variable. How this plays out exactly will become clear once my formal account has been presented.

I now give a brief outline of the syntax and semantics of a basic formal language L_0 for subjunctive conditionals. Suppose given a basic stock of predicate symbols, as in standard first-order logic. We construe these predicate symbols as being in indicative mood, and add, for each such symbol P, a subjunctive predicate symbol P^s of the same arity. Predicate symbols of either type, indicative or subjunctive, can be prefixed to sequences of singular terms of appropriate length to produce the atomic formulas of L_0.

The familiar truth-functional connectives and the first-order quantifiers can be used to generate new formulas from old, as can the conditional operator $\Rightarrow$, with the (temporary) proviso that both antecedent and consequent of a $\Rightarrow$-conditional must contain at least one subjunctive predicate.[13] For the purposes of the present section, I disallow nested occurrences of $\Rightarrow$, so that, in an L_0-formula $A \Rightarrow B$, neither A nor B are themselves permitted to contain occurrences of $\Rightarrow$. The no-nesting constraint will be relaxed, though not completely abolished, in §4.[14]

With a tip o' the hat to Dudman, we call an L_0-formula an L_0-*sentence* if it contains no free occurrences of individual variables (i.e. if it is *individually closed*) and no occurrences of the subjunctive marker that do not lie within the scope (i.e. antecedent or consequent) of a double arrow (i.e. if it is *subjunctively closed*). As we will see, L_0-sentences correspond to those sentences of English that are truth-evaluable *tout court* (i.e. without having to specify additional parameters, like counterfactual scenarios or assignments of values to pronouns).

We next define what it means for a world w to satisfy an L_0-formula A with respect to a variable assignment σ (briefly: for w to σ-satisfy A). To make this precise, I will formulate the definitions in terms of model-theoretic stipulations for possible-worlds structures. Such structures **M** consist of a set W of possible worlds, whose distinguished element @(**M**) is called the actual world of **M**; a non-empty set D of individuals (for

[13] This proviso reflects the facts that (a) English grammar does not countenance conditionals with purely indicative antecedents (consequents) and partially or fully subjunctive consequents (antecedents), and (b) I will only suggest provisions for the case where both antecedent and consequent are purely indicative in §4.
[14] It is possible to be much more generous in allowing nested occurrences of the double arrow than I am here or in §4, but space considerations prevent me from addressing the issue in this paper.

simplicity, we assume that the same individuals exist in all worlds); and, for each *n*-ary indicative predicate symbol P and each world *w* in W, an extension P_w of P with respect to *w*. We will write '$\mathbf{M} \vDash_w A\ [\sigma]$' for '*w* σ-satisfies A in the structure **M**'.

A world *w* σ-satisfies a *subjunctive* atomic formula $R^s xy$, say, just in case the pair of individuals $<\sigma(x), \sigma(y)>$ is a member of the extension R_w of R at *w*. Invoking the extension of R at *w*, rather than its extension at the actual world, is apposite because the subjunctive mood signals that we are making reference to whatever possible world is the salient counterfactual scenario (i.e. here, by stipulation, the world *w*). Formally we define: $\mathbf{M} \vDash_w R^s xy\ [\sigma]$ if and only if $<\sigma(x), \sigma(y)> \in R_w$.

By contrast, *w* σ-satisfies an *indicative* atomic formula R*xy* if and only if $<\sigma(x), \sigma(y)>$ is a member of the extension $R_{@(\mathbf{M})}$ of R at the actual world @(**M**). This is plausible because, in the context of counterfactual discourse, the indicative is used to refer to how things really are, regardless of what counterfactual scenario may be salient. Model-theoretically: $\mathbf{M} \vDash_w Rxy\ [\sigma]$ if and only if $<\sigma(x), \sigma(y)> \in R_{@(\mathbf{M})}$.

The clauses for the propositional connectives are standard; e.g. *w* σ-satisfies ~A if and only if it is not the case that *w* σ-satisfies A, and *w* σ-satisfies A⊃B just in case A is not σ-satisfied by *w*, or B is σ-satisfied by *w*. In model-theoretic terms: $\mathbf{M} \vDash_w {\sim}A\ [\sigma]$ if and only if not $\mathbf{M} \models_w A\ [\sigma]$; and $\mathbf{M} \models_w A \supset B\ [\sigma]$ if and only if either not $\mathbf{M} \models_w A\ [\sigma]$ or $\mathbf{M} \models_w B\ [\sigma]$.

The quantifier clause, too, is as expected: for *w* to σ-satisfy $\exists x$A is for there to be an *x*-variant τ of σ such that A is τ-satisfied by *w*. Formally: $\mathbf{M} \models_w \exists x A\ [\sigma]$ if and only if there is some *x*-variant τ of σ such that $\mathbf{M} \models_w A\ [\tau]$.

We finally turn to the satisfaction conditions for formulas of the form A⇒B. Since it simplifies exposition considerably, I will avail myself of a version of Lewis's Limit Assumption[15] and stipulate that *w* σ-satisfies A⇒B if and only if B is σ-satisfied by all worlds *v* closest[16] to @(**M**) that σ-satisfy A. Expressed model-theoretically, this means that $\mathbf{M} \models_w A \Rightarrow B\ [\sigma]$ if and only if $\mathbf{M} \models_v B\ [\sigma]$ for all worlds *v* with the properties that (a) $\mathbf{M} \models_v A\ [\sigma]$ and (b) for no world *u* with $u <_{@(\mathbf{M})} v$, $\mathbf{M} \models_u A\ [\sigma]$.

An individually closed L_0-formula is σ-satisfied by *w* for some σ just in case it is σ-satisfied by *w* for any σ. For individually closed formulas, the notion of being satisfied by a world is thus independent of the choice of a variable assignment, and we can define

---

[15] To be precise, we assume that whenever there is some world *v* that σ-satisfies A, then there are worlds *u* with the following two properties: (i) *u* σ-satisfies A, and (ii) whenever any world *w* σ-satisfies A, then *u* is at least as close to @(**M**) as *w* is. That is, if A is σ-satisfied anywhere, there are worlds *u* closest to @(**M**) that σ-satisfy A.

[16] We assume that the notion of closeness of worlds is governed by something like Lewis's (1973, 48) clauses (1) through (6) for comparative similarity. For the purposes of this section, we need only worry about closeness to the actual world, due to our prohibition against nested occurrences of '⇒'.

such a formula to be *satisfied by w* just in case it is σ-satisfied by *w* for all σ (equivalently, some σ). Speaking model-theoretically, if A is individually closed, we have that for any σ and τ, $\mathbf{M} \models_w A\ [\sigma]$ if and only if $\mathbf{M} \models_w A\ [\tau]$. Thus in this case we define $\mathbf{M} \models_w A$ as $\mathbf{M} \models_w A\ [\sigma]$ for all (equivalently, some) σ.

Similarly, there is some *w* that σ-satisfies a subjunctively closed L_0-formula if and only if all *w* σ-satisfy it; that is, a subjunctively closed formula's being σ-satisfied is independent of the choice of a world of evaluation. We thus define a subjunctively closed formula to be σ-*satisfied* just in case it is σ-satisfied by all *w* (equivalently, some *w*). Formally: If A is subjunctively closed, we have for all *v* and *w* that $\mathbf{M} \models_v A\ [\sigma]$ if and only if $\mathbf{M} \models_w A\ [\sigma]$. Thus we define $\mathbf{M} \models A\ [\sigma]$ to mean $\mathbf{M} \models_v A\ [\sigma]$ for every (equivalently, some) *v*.

Finally, an L_0-*sentence* is σ-satisfied by some *w* for some σ if and only if it is σ-satisfied by any *w* for any σ, so for L_0-sentences being satisfied is independent of the choice of a world of evaluation and of a variable assignment. We thus define: An L_0-sentence A is *true* if and only if it is σ-satisfied by some *w* for some σ (equivalently, by all *w* for all σ). Expressed model-theoretically, this means that if A is an L_0-sentence, we have for all *v*, *w*, σ, and τ: $\mathbf{M} \models_v A\ [\sigma]$ if and only if $\mathbf{M} \models_w A\ [\tau]$. Thus we may define $\mathbf{M} \models A$ to mean $\mathbf{M} \models_v A\ [\sigma]$ for all (equivalently, some) *v* and σ.

The relationship of our language to Lewis's is as follows: Given formulas A and B of ordinary predicate logic, Lewis's subjunctive conditional A > B has the same truth conditions as the L_0-formula $A^s \Rightarrow B^s$, where A^s arises from A by attaching the superscript 's' to every predicate that occurs in A, and similarly for B and B^s. In other words, for purely subjunctive conditionals, our account predicts the same truth conditions as Lewis's. But note that not all L_0-conditionals of the form $C \Rightarrow D$ can be obtained in this way from Lewis conditionals, because C or D may contain indicative predicates (which A^s and B^s cannot).

Our semantics handles subjunctive-indicative conditionals easily. If we apply it to our proposed formalization (3s) of (3), we obtain the following truth condition:

(3s′) For all worlds *u* closest to the actual world @ with the property that everybody who voted for Christa in @ voted for Barbara in *u*, it is the case that Anna wasn't elected chair in *u*.

That delivers the right result: Our scenario suggests that the appropriate measure of closeness to @ is the number of faculty who vote just as they do in @. The worlds *u* closest to @ that satisfy the antecedent $\forall x\,(Vxc \supset V^s xb)$ of (3s) must then have everybody except the Christa-supporters-in-@ voting in *u* as they did in @. This means that in any such world *u*, Barbara garners six votes and Anna five, so that Anna isn't elected chair. Thus (3s′) is indeed the right truth condition for (3).

Similarly, the truth condition our semantics delivers for the formalization (4s) of (4) is:

(4s′) For all worlds *u* closest to @ with the property that Barbara was elected chair in *u*, it is the case that some books the student is checking out in @ were thrown out in *u*.

This, too, is correct. Closeness to the actual world is to be cashed out, it may be assumed, in terms of Barbara's faithfulness to the promises she made while campaigning in @. The worlds *u* closest to @ such that Barbara was elected chair in *u* have her carrying out the announced plan of getting rid of some of the logic books. In view of the fact that the logic books are just the books the student is checking out, it follows that all such worlds *u* satisfy the consequent $\exists y\ (\mathrm{C}y\ \&\ \mathrm{T}^{\mathrm{s}}y)$ of (4s). Thus (4s′) is indeed the right truth condition for (4).

In general, to formalize an English sentence in L_0, one translates English indicative predicates into indicative predicate symbols of L_0, English subjunctive predicates into subjunctive predicate symbols of L_0, 'if' into '$\Rightarrow$', and otherwise proceeds just as for ordinary predicate logic. We noted above, conceding Dudman's point, that antecedents and consequents of subjunctive conditionals, such as 'Shakespeare had not written *Hamlet*' and 'Somebody other than Shakespeare would have written *Hamlet*', are not complete English sentences. This is due, it would seem, to their containing subjunctive predicates that are not governed by a conditional operator. Consequently, all formalizations of *complete* English sentences are subjunctively closed L_0-formulas (indeed, L_0-sentences). And this is as it should be, for the subjunctively closed L_0-formulas are precisely those that possess a truth-value independently of a prior choice of a world of evaluation—just as complete English sentences have their truth-values independently of any counterfactual circumstances with respect to which subjunctive predicates could be evaluated. This will become relevant in §5 below.

4. Indicative Conditionals

The account I have given predicts the correct truth conditions for purely subjunctive and for subjunctive-indicative conditionals, and it explains, in an intuitively satisfying way, the interaction of *if* with the grammatical moods in such conditionals. In order to fulfill desiderata (A)-(C), however, it will be necessary to extend the account to indicative conditionals as well. To this task I now turn.

Before we begin, I must address a worry that arises from the framework as developed so far. If we are to use the single operator ⇒ for subjunctive *and* indicative conditionals, we are committed to regarding indicative conditionals as being of the form A⇒B, with purely indicative A and B. But if indicative predicates are invariably evaluated at the actual world, simple extrapolation of the semantic clause governing ⇒ in L_0 yields the following for such an indicative conditional:[17]

(*) A⇒B is true just in case all worlds *w* closest to @ that satisfy A also satisfy B.

[17] For simplicity, we assume that A and B are individually closed and thus make do without variable assignments.

Since A and B are subjunctively closed (being purely indicative, they contain no subjunctives whatsoever), satisfaction of A and B is independent of the world of evaluation; that is, for any worlds *v* and *w*, *v* satisfies A just in case *w* does, and similarly for B. Recall that, by definition, A is true if and only if all worlds satisfy it (equivalently, if some world satisfies it). Suppose first that A is indeed true. Then every world satisfies A, and in particular @, the unique world closest to @, does. By (*), this means A⇒B is true just in case @ satisfies B which, by subjunctive closure of B, is equivalent to B's being true. In other words: When A is true, the truth-value of A⇒B is identical to the truth-value of B. But if A is not true, by subjunctive closure of A, no world satisfies A, so (*) is vacuously true, and hence A⇒B is true. Taken together, this shows that A⇒B is the material conditional: true if A is false, and identical in truth-value to B in case A is true. If indicative conditionals are not material conditionals, this is unacceptable.

We thus face a challenge: Is it possible to maintain that indicative predicates must always be interpreted in the actual world, without thereby making the indicative conditional collapse into the material one? The answer is, fortunately, yes.

Weatherson (2001) has argued that, while the metaphysical feel of subjunctive conditionals comes from their quantifying over worlds of evaluation (i.e. worlds considered as counterfactual), the epistemic feel of indicative conditionals derives from the fact that they vary instead the world considered as actual. We appropriate this insight for our framework in the following way: If in a conditional A⇒B both A and B contain free subjunctives, the truth conditions stated in §3 apply; in particular, the conditional operator acts as a quantifier over worlds of evaluation, and we are dealing with a 'metaphysical' (subjunctive) conditional. But when neither A nor B contain free subjunctives, so that the conditional operator would run empty when acting as a quantifier over worlds of evaluation, it instead lets vary the world considered as actual, and we are dealing with an epistemic (indicative) conditional.[18,19] Here, then, is the formal account.

The language L_1 has the same vocabulary as the language L_0 of §3, in particular, it contains, with each indicative predicate symbol P, also its subjunctive version P^s. Atomic formulas are defined as in L_0, and new formulas can be constructed from old using the truth-functional connectives and the first-order quantifiers as before. The clause for the double arrow is modified to read as follows: If A and B are L_1-formulas, then so is $A \Rightarrow B$, provided that (a) neither A nor B contains an occurrence of the double arrow, but both contain at least one occurrence of the subjunctive marker (in which case we call $A \Rightarrow B$ an s-conditional), or (b) both A and B are subjunctively closed, and whenever A or B contains a subformula of the form $C \Rightarrow D$, this subformula is an s-conditional (in which

[18] This may turn out to be an attractive analysis for all broadly modal operators that do not subjunctivize predicates within their scopes. For instance, regarding the sentence 'Peter believes that Mary is in Boston', one might balk at the claim that the indicative 'is' pertains to the actual world. However, analyzing the sentence along the lines of 'Every world *w* compatible with Peter's beliefs is such that, if *w* is considered actual, "Mary is in Boston" is true (i.e. true at *w*)' makes this much more plausible. Contrast this with 'Peter wishes that Mary were in Boston'; here Peter's attitude is clearly directed toward worlds he considers as counterfactual, and consequently, the subjunctive mood is required.

[19] Our formal language L_1 introduced below allows for a limiting case of conditional in which the antecedent and consequent are subjunctively closed *and* fail to contain an indicative predicate. The truth conditions of such conditionals turn out to be those of material conditionals, which is as it should be. English counterparts are statements of doubtful grammaticality like 'If (Oswald wouldn't have shot Kennedy had he overslept the morning of the assassination), then (if Kennedy had served out his term, Johnson wouldn't have become President)'. I suspect that, if one grants that these are well-formed sentences, there is no option, on any account, other than treating them as material conditionals.

case we call $A \Rightarrow B$ an i-conditional).[20] As before, an L_1-formula is subjunctively closed if it contains no subjunctive marker not within the scope of a double arrow, and individually closed if it contains no free occurrences of individual variables. An L_1-sentence is an L_1-formula that is both subjunctively and individually closed.

Before we turn to the semantics for L_1, let us consider some examples of i- and s-conditionals. We've already seen

(3s) $[\forall x(\mathrm{V}xc \supset \mathrm{V}^{s}xb)] \Rightarrow \sim \mathrm{E}^{s}a$

and

(4s) $\mathrm{E}^{s}b \Rightarrow \exists y(\mathrm{C}y \;\&\; \mathrm{T}^{s}y)$.

Both of these are s-conditionals, because their protases and apodoses all contain free subjunctives. Standard purely indicative conditionals such as (2) correspond to i-conditionals, in this case to $\neg \mathrm{H}s \Rightarrow \exists x(x \neq s \;\&\; \mathrm{H}s)$, where s represents 'Shakespeare' and

[20] The prefixes 'i' and 's' stand for 'indicative' and 'subjunctive', respectively. The restriction on the formation of i-conditionals—that their antecedents and consequents must not themselves contain i-conditionals—reflects the fact that indicative conditionals do not seem naturally to occur nested. Weatherson (2001) tentatively suggests that sometimes we can make sense of such embeddings, but it seems to me his example does not establish the point. On the other hand, as we will see below, indicative conditionals do embed subjunctive ones, which L_1 permits. One may want to allow *some* ways of nesting i-conditionals, like $A \Rightarrow (B \Rightarrow C)$, since their acceptability seems to be conceded by some commentators. The semantics of such right-nested i-conditionals could be accommodated, but I will refrain from expanding on the issue here. Similarly, we could be more generous in terms of nesting s-conditionals, as mentioned in §3, although I find it doubtful that nested subjunctive conditionals really are grammatical in English. Note also that s-conditionals cannot embed i-conditionals, which is consistent with the fact that English subjunctive conditionals do not seem to permit embedding of indicative conditionals.

H*x* stands for ‘*x* wrote *Hamlet*’. But there are more interesting examples. Consider the following, due to John Burgess:

(5) If everybody who, had they not been paid to show up, wouldn’t have come, left by seven o’clock last night, Wilbur was left without party guests.

This sentence formalizes in L_1 as

(5s) $\forall x[(\sim P^s x \Rightarrow \sim C^s x) \supset Lx] \Rightarrow \sim \exists x Wx$,

which is an i-conditional because its antecedent and consequent are subjunctively closed, and the ⇒-subformula of the antecedent is an s-conditional.[21]

Here is a similar example, which has the s-conditional embedded in the consequent rather than the antecedent.

(6) If Paul attended the party, then someone who would have had a good time if Paul had not come got a big surprise.

In the language L_1, we formalize it as the following i-conditional:

[21] In binary quantifier notation, (5s) becomes $([\forall x\colon \sim P^s x \Rightarrow \sim C^s x]\ Lx) \Rightarrow \sim \exists x Wx$.

(6s) $\mathrm{A}p \Rightarrow \exists x[(\sim\mathrm{A}^{\mathrm{s}}p \Rightarrow \mathrm{G}^{\mathrm{s}}x)\ \&\ \mathrm{S}x]$.

We now define the semantics for L_1. This will be done by relativizing the semantics for L_0 to a contextual parameter, called an *epistemic base*, and by providing an additional clause for the evaluation of i-conditionals (which is the only case in which the epistemic base comes into play). For simplicity, we define an epistemic base β to be a non-empty set of possible worlds. The elements of β should be thought of as the possible worlds that may, for all that is known, be actual, or in other words, the worlds compatible with the knowledge about the actual world available in the context of utterance.[22]

We begin by defining, for a given variable assignment σ, the notion of a world *w*'s σ-satisfying an L_1-formula A relative to the epistemic base β in a given possible-worlds structure **M**, in symbols $\mathbf{M} \models^{\beta}_{w} \mathrm{A}\ [\sigma]$. As usual, we assume that @(**M**) is the distinguished actual world of **M**. For any world *v*, we let $\mathbf{M}_v$ be the possible-worlds structure that is just like **M** except that its distinguished actual world, $@(\mathbf{M}_v)$, is *v*.

A possible world σ-satisfies R*xy* relative to β just in case <σ(*x*), σ(*y*)> belongs to the extension $\mathrm{R}_{@(\mathbf{M})}$ of R at the actual world @(**M**) of the structure **M**: $\mathbf{M} \models^{\beta}_{w} \mathrm{R}xy\ [\sigma]$ if and only if $\langle \sigma(x), \sigma(y) \rangle \in \mathrm{R}_{@(\mathbf{M})}$.

[22] The epistemic base could itself be given by means of a set S of L_0-sentences such that a world *w* belongs to β just in case every member of S would be true (in the sense of §3) if *w* were actual. Instead of thinking of the epistemic base as the set of worlds compatible with what is known, one might take it to be the set of worlds consistent with the presuppositions in force in the context of utterance, or other broadly epistemic notions. There is also the question whose knowledge (presuppositions, etc.) determines β. Cf. (von Fintel and Gillies 2008).

A world w σ-satisfies $R^s xy$ relative to β just in case the pair <σ(x), σ(y)> belongs to the extension R_w of R at w: $\mathbf{M} \models^{\beta}_{w} R^s xy\ [\sigma]$ if and only if $\langle\sigma(x), \sigma(y)\rangle \in R_w$.

A world σ-satisfies A & B relative to β just in case it σ-satisfies both A and B relative to β, and similarly for the other truth-functional connectives: $\mathbf{M} \models^{\beta}_{w} A \;\&\; B\ [\sigma]$ if and only if $\mathbf{M} \models^{\beta}_{w} A\ [\sigma]$ and $\mathbf{M} \models^{\beta}_{w} B\ [\sigma]$, etc.

A world σ-satisfies $\exists x$A relative to β just in case there is an x-variant τ of σ such that the world in question τ-satisfies A relative to β: $\mathbf{M} \models^{\beta}_{w} \exists x A\ [\sigma]$ if and only if there is an x-variant τ of σ such that $\mathbf{M} \models^{\beta}_{w} A\ [\tau]$.

Let A⇒B be an s-conditional. Then w σ-satisfies A⇒B relative to β just in case the worlds v closest to the actual world that σ-satisfy A relative to β also σ-satisfy B relative to β: $\mathbf{M} \models^{\beta}_{w} A \Rightarrow B\ [\sigma]$ if and only $\mathbf{M} \models^{\beta}_{v} B\ [\sigma]$ for all worlds v with the two properties that (a) $\mathbf{M} \models^{\beta}_{v} A\ [\sigma]$ and (b) for no u with $u <_{@(\mathbf{M})} v$, $\mathbf{M} \models^{\beta}_{u} A\ [\sigma]$.

Note that up to this point the epistemic base has been idling, and therefore, when A is an L_0-formula, $\mathbf{M} \models^{\beta}_{v} A\ [\sigma]$ holds if and only if $\mathbf{M} \models_{v} A\ [\sigma]$ in the sense of §3, regardless of choice of β. If, moreover, A is a subjunctively closed L_0-formula (as will be the case with antecedents and consequents of i-conditionals), whether or not $\mathbf{M} \models_{v} A\ [\sigma]$ is independent of *v*, so $\mathbf{M} \models^{\beta}_{v} A\ [\sigma]$ if and only if $\mathbf{M} \models A\ [\sigma]$. This should be kept in mind while considering the following satisfaction clause for i-conditionals.

If A⇒B is an i-conditional, *w* σ-satisfies A⇒B relative to β just in case all worlds *v* in β with the property that A would be σ-satisfied if *v* were actual are such that B would be σ-satisfied if *v* were actual: $\mathbf{M} \models^{\beta}_{w} A \Rightarrow B\ [\sigma]$ if and only if for all $v \in \beta$, if $\mathbf{M}_v \models A\ [\sigma]$, then $\mathbf{M}_v \models B\ [\sigma]$.[23]

Satisfaction of L_1-*sentences* is independent of variable assignments and worlds of evaluation (though not, in general, of epistemic bases). We can thus define an L_1-sentence to be true relative to a base β just in case for all worlds *w* and all variable assignments σ, *w* σ-satisfies it relative to β (equivalently, some world *w* σ-satisfies it relative to β, for some σ). Again, in model-theoretic terms we have, for L_1-sentences A, worlds *v* and *w*, and assignments σ and τ, that $\mathbf{M} \models^{\beta}_{w} A\ [\sigma]$ if and only if $\mathbf{M} \models^{\beta}_{v} A\ [\tau]$, and hence we may define $\mathbf{M} \models^{\beta} A$ to mean $\mathbf{M} \models^{\beta}_{w} A\ [\sigma]$ for all (equivalently, some) *w* and σ.

[23] Recall that $\mathbf{M}_v$ is just like $\mathbf{M}$ except that the distinguished actual world of $\mathbf{M}_v$ is *v*. Some modifications to our stipulation would need to be made if we countenanced nested indicative conditionals like A⇒(B⇒C). Note also that, since evaluation of i-conditionals changes which world is considered as actual, and i-conditionals may embed s-conditionals, we now need a notion of closeness between arbitrary worlds, and not just between any given world and the *actual* actual world.

As an application of the semantics for L_1, let us consider our earlier example

(6) If Paul attended the party, then someone who would have had a good time had Paul not come, got a big surprise

and its proposed formalization

(6s) $Ap \Rightarrow \exists x\, [(\sim A^s p \Rightarrow G^s x)\ \&\ Sx]$.

According to our semantics, (6s) is true relative to β just in case for all w and σ, w σ-satisfies (6s) relative to β.

So let σ and w be arbitrary. By definition, w σ-satisfies (6s) relative to β just in case all worlds v in β are such that, if (i) Ap is σ-satisfied provided v is actual, then, provided v is actual, (ii) $\exists x\, [(\sim A^s p \Rightarrow G^s x)\ \&\ Sx]$ is also σ-satisfied. Since Ap is indicative, condition (i) means that the interpretation of p belongs to the extension A_v of A at v (note that when v is considered actual, indicative predicates are evaluated with respect to v). Condition (ii) means that for some individual q belonging to the extension S_v of S at v (again, since v is considered actual, the indicative predicate S is evaluated with respect to v), the formula $(\sim A^s p \Rightarrow G^s x)$ is τ-satisfied by the x-variant τ of σ that maps x to q, provided v is actual. Thus condition (ii) holds if and only if there is an individual q in S_v such that for all worlds u closest to v with the property that p does not belong to A_u, q belongs to G_u.

Putting things together, this means that (6s) is true just in case for all $v \in \beta$, if $p \in A_v$, then there is some $q \in S_v$ such that $q \in G_u$ for all worlds u closest to v with the property that $p \notin A_u$. In other words, our semantics predicts for (6) to be true if and only if there exists, for every world v compatible with what is known and in which Paul attended the party, an individual who is surprised in v and who has a good time in every world u that is maximally close to v with the property that Paul didn't attend the party in it. I submit that this is just the right truth condition for (6).[24]

5. Actuality

The actuality operator **A** is not a part of Stalnaker's or Lewis's official accounts, but lends itself to inclusion in either.[25] With its help, the Stalnaker-Lewis framework admits of the following formalization of (3):

(3**A**) $\forall x\, (\mathbf{A}Vxc \supset Vxb) > {\sim}Ea$,

and similarly for (4):

(4**A**) $Eb > \exists y\, (\mathbf{A}Cy \,\&\, Ty)$.

[24] Similarly, one can show that (2) is true relative to a base β just in case for every world u in β, if Shakespeare didn't write *Hamlet* in u, then there is an individual other than Shakespeare who wrote *Hamlet* in u. This is, of course, a standard account of the truth conditions of purely indicative conditionals.

[25] See e.g. (Crossley and Humberstone 1977). A formula **A**C is true at a world w (under a given variable assignment) just in case C is true at the actual world (under the same variable assignment). Here truth at a world is of course defined in the standard Stalnaker-Lewis way.

It is easy to check that (3**A**) and (4**A**) do indeed have the same truth conditions, under standard Stalnaker-Lewis semantics, as our (3s) and (4s) do under the semantics proposed in §3.

Given the availability of this simple fix to extant theories, why even consider the analysis proposed above? Before I can answer this challenge, it is necessary to probe a little deeper into the workings of the actuality operator. Consider

> (7) Oswald shot Kennedy, but if he hadn't shot him, somebody else would have.

The Stalnaker-Lewis language with actuality provides two options for formalizing (7), to wit:

> (7a) $\mathrm{S}o\ \&\ [\neg \mathrm{S}o > \exists x\ (x \neq o\ \&\ \mathrm{S}x)]$

> (7b) $\mathbf{A}\mathrm{S}o\ \&\ [\neg \mathrm{S}o > \exists x\ (x \neq o\ \&\ \mathrm{S}x)]$

In the latter case, the fact that the first conjunct in (7) speaks about the actual world is guaranteed by including it within the scope of **A**. In the former, the same effect is achieved by the convention that, in the Stalnaker-Lewis framework, truth *tout court* is understood to be truth at the actual world. Since 'S*o*' does not occur within the scope of any operator that might change the world of evaluation, the actual world, *qua* default world, is used for the evaluation of 'S*o*'.

Now if both 'S*o*' and '**A**S*o*' are permissible formalizations of 'Oswald shot Kennedy', we face the oddity that the modal profile of this English sentence (and in fact *any* indicative sentence) is underdetermined, for 'S*o*' is contingent, whereas '**A**S*o*' is either necessarily true (if 'S*o*' is true at the actual world) or necessarily false (if 'S*o*' is false at the actual world). Moreover, there seem to be no contextual features that could resolve the ambiguity. It seems rather implausible that our quotidian, matter-of-fact statements should be subject to such radical ambiguity. What is more, as far as I can see there is no evidence whatsoever for the two kinds of indicative that this approach predicts —a standard indicative useable only outside modal constructions, and a 'superindicative', distinguished by the presence of **A**, which can be used both inside and outside modal constructions and invariably produces non-contingent utterances. Arguably, then, this approach is in conflict with our desideratum (B), that is, the requirement that a theory of conditionals illuminate the interaction of *if* with the indicative and subjunctive moods, for it postulates a non-observable distinction between two purported indicative moods.

There are two ways in which this ambiguity could be excised from the actuality-enriched Stalnaker-Lewis language: First, by insisting that, in formalizing English declarative sentences, indicative verbs must always be rendered as predicate symbols governed by **A**, thereby excluding (7a) as an admissible formalization of (7). This, in effect, turns the Stalnaker-Lewis language with actuality into a mere notational variant of (the subjunctively closed fragment of) our language L_0: presence of **A** corresponds to absence of the subjunctive marker and vice versa. In other words, instead of marking the subjunctive and leaving the indicative unmarked, as we have done, such a language marks the indicative and leaves the subjunctive unmarked. Our notation seems preferable because ordinary non-modal predicate logic is literally embedded in it as a sublanguage. For instance, we express ‘Oswald shot Kennedy’ in ordinary predicate logic as ‘S*o*’, and that is exactly the formalization this sentence receives in L_0. The variant of the Stalnaker-Lewis language under discussion, by contrast, must insist that when we move from non-modal predicate logic to the language of conditionals, sentences must be prefixed with an additional ‘**A**’, that is, the first-order formalization ‘S*o*’ of ‘Oswald shot Kennedy’ must first be explicitly indicativized to ‘**A**S*o*’ before it becomes part of the actuality-enriched Stalnaker-Lewis language.

The second way of removing the ambiguity in modal profile consists in stipulating that indicative predicates are formalized without any marking when they occur outside all modal scopes, but must be prefixed with an actuality operator when they occur within the antecedent or consequent of a subjunctive conditional (or within the scope of another metaphysically modal operator). This would exclude (7b) as a formalization of (7). The ad hoc character of the move, however, is readily apparent: Why stipulate two logical forms, ‘S*o*’ and ‘**A**S*o*’, for the single indicative statement ‘Oswald killed Kennedy’, depending on its context of occurrence? Moreover, why stipulate that one and the same logical form, ‘S*o*’, can represent an indicative predicate when not within the scope of a conditional operator (as in the first conjunct of (7a)) and a subjunctive predicate when occurring within such a scope (as in the antecedent of (7a))?[26] There does not seem to be any empirical evidence to support such a stipulation regarding logical forms. Just as in the original, ambiguity-encumbered language with actuality, we have here not an illumination, but a confusion of the role of the indicative mood in relation to subjunctive conditionals, and thus, arguably, a violation of desideratum (B).

If these observations regarding the deployment of **A** are correct, it follows that the only unobjectionable implementation of the actuality operator in the Stalnaker-Lewis language amounts to a mere notational variant of the language L_0 proposed in §3.

6. Conclusion

How well does the account of conditionals developed in sections 3 and 4 satisfy the desiderata laid down at the beginning of this essay? To recapitulate: We suggested that a satisfactory logical analysis of English *if*-sentences should

> (A) predict correct truth conditions for as wide a class of English conditionals as possible,
>
> (B) illuminate the interaction of *if* with the indicative and subjunctive moods in all grammatically admissible configurations, and
>
> (C) explain how the single English word *if* is able to perform the ostensibly distinct tasks it does in indicative and subjunctive conditionals.

[26] Related objections are raised, in the context of modal logic, by Wehmeier (2005).

I believe that the account I have outlined in sections 3 and 4, when compared with extant theories, fares well with respect to all three criteria.

As regards (A), we have seen that Stalnaker's and Lewis's approaches are unable to accommodate subjunctive-indicative conditionals like (3) and (4), whereas our account takes them in stride. On other conditionals—purely subjunctive and indicative—our account agrees with standard analyses. Overall, we have therefore extended the domain of English conditionals for which correct truth conditions can be predicted.[27]

With respect to (B), our account shows that *if* binds the subjunctive mood much like a quantifier binds a world variable, but treats indicative mood as a constant (for the actual world). In the absence of bindable subjunctives, *if* escapes triviality (i.e. collapse into the material conditional) by varying which world is being considered as actual. This two-dimensional model elucidates the interaction of *if* with mood in purely subjunctive, subjunctive-indicative, and indicative conditionals, and it furthermore explains the epistemic flavor of indicative conditionals.

Finally, regarding (C), we have been able to explain the difference in semantic rôle of *if*—epistemic versus metaphysical—by reference to the presence or absence of subjunctive binding, without having to posit lexical ambiguity like Lewis. If both antecedent and consequent of an *if* sentence are subjunctively closed, we have an indicative conditional: There are no subjunctive world variables for the conditional

[27] Augmenting the Stalnaker-Lewis framework with the actuality operator removes its expressive deficit, but, as argued in §5, at the cost of either turning the framework into an awkward notational variant of ours, or of running seriously afoul of desideratum (B).

operator to bind, therefore it turns to the other dimension of variability and changes the world considered as actual (within the constraints of the epistemic base). If, on the other hand, both antecedent and consequent contain free subjunctives, we have a subjunctive conditional: The conditional operator interprets the subjunctive as a variable over worlds of evaluation (but leaves fixed the world considered as actual, which is called upon for the evaluation of indicative predicates, if any).

This way of distinguishing between indicative and subjunctive conditionals would come under pressure if there were *if* sentences whose antecedents (consequents) contained free subjunctives while their consequents (antecedents) were subjunctively closed. As far as I can tell, however, no such *if* sentences are grammatical in English, which would seem to corroborate the account proposed here. Similarly, if a subjunctive conditional A⇒B could embed an indicative conditional C⇒D, the question would arise whether the interpretation of an indicative predicate occurring in, say, C must remain fixed (as suggested by our semantics for the s-conditional A⇒B) or is allowed to vary (as suggested by our semantics for the i-conditional C⇒D). Again, embedding indicative conditionals within subjunctive ones does not appear to be grammatical in English, which I take to further support my view.

The account advanced in this paper both supports and is supported by Wehmeier's (2004) analysis of metaphysically modal statements by means of his formal language SML containing explicitly subjunctivized predicates. This is significant because the semantic picture that emerges is quite different from the received, Kripke-inspired view. For if grammatical mood attaches directly to predicates, rather than being a scopal property of complete sentences, there seems to be no basis for the claim that definite descriptions are typically non-rigid designators. After all, it is a consequence of the type of framework espoused in this paper that an indicative description, say "the man who wrote *Hamlet*", always evaluates to the individual who wrote *Hamlet* in the actual world, no matter where the description occurs in a sentence, and is thus a rigid designator. Its subjunctive cousin, "the man who would have written *Hamlet*", is non-rigid all right, but arguably fails to be a designator, for it is impossible to say which individual is its referent without prior specification of a salient counterfactual scenario. If the picture I have painted here of subjunctive and indicative conditionals, of their logical form and semantics, is correct, this should thus have ramifications for the theory of reference and philosophy of language more generally.[28]

Bibliography.

Stephen Barker, 'Parsing *if*-sentences and the conditions of sentencehood,' *Analysis* 56 (4), 1996, 210-218.

John Burgess, 'Counterfactuals and Subjunctives,' unpublished ms, 2011.

[28] I have recently (November 2010) seen an unpublished paper by John Mackay titled 'Conditionals, Mood and Binding' that reaches some of the same conclusions, and proposes similar revisions to the orthodoxy, as I do in this essay. It seems fair to say that Mackay's focus is on the linguistic side, while mine is logic oriented. In any case, it is to be hoped that his highly relevant and interesting work will soon be available in published form.

Crossley, J.N. and Lloyd Humberstone, ‘The Logic of “Actually”’, *Reports on Mathematical Logic* 8, 1977, 11-29.

Davies, Martin and Lloyd Humberstone, “Two Notions of Necessity”, *Philosophical Studies* 38, 1980, 1-30.

Victor H. Dudman, ‘Parsing ‘If’ Sentences,’ *Analysis* 44 (4), 1984, 145-153.

Gareth Evans, ‘Reference and Contingency,’ *The Monist* 62 (2), 1979, 178-213.

Kai von Fintel and Anthony S. Gillies, ‘CIA Leaks,’ *Philosophical Review* 117 (1), 2008, 77-98.

Lloyd Humberstone, ‘Scope and Subjunctivity,’ *Philosophia* 12, 1982, 99-126.

Lloyd Humberstone, ‘Two-Dimensional Adventures,’ *Philosophical Studies* 118, 2004, 17-65.

Sabine Iatridou, ‘The grammatical ingredients of counterfactuality,’ *Linguistic Inquiry* 31 (2), 2000, 231–270.

David Lewis, *Counterfactuals*, Blackwell 1973.

David Lewis, ‘Probabilities of Conditionals and Conditional Probabilities’, *Philosophical Review* 85, 1976, 297-315.

John Mackay, ‘Conditionals, Mood and Binding’, unpublished manuscript.

Donald Nute and Charles B. Cross, ‘Conditional Logic,’ in D.M. Gabbay and F. Guenthner, *Handbook of Philosophical Logic*, 2nd edition, volume 4, Kluwer 2002, 1-98.

Orin Percus, ‘Constraints on some other variables in syntax,’ *Natural Language Semantics* 8 (3), 2000, 173–229.

Robert Stalnaker, ‘A Theory of Conditionals,’ in N. Rescher (ed.), *Studies in Logical Theory*, Blackwell, 1968, 98-112.

Robert Stalnaker, ‘Indicative conditionals,’ *Philosophia* 5 (3), 1975, 269-286.

Jordan Stein, *How Many Notions of Necessity?*, doctoral dissertation, University of California, Irvine, 2011.

Brian Weatherson, ‘Indicatives and Subjunctives,’ *Philosophical Quarterly* 51, 2001, 200-216.

Kai F. Wehmeier, ‘In the Mood,’ *J. Philos. Logic* 33 (6), 2004, 607-630.

Kai F. Wehmeier, ‘Modality, Mood, and Descriptions,’ in R. Kahle (ed.), *Intensionality: An Interdisciplinary Discussion*, AK Peters, 2005, 187-216.

Kai F. Wehmeier, ‘Subjunctivity and Cross-World Predication,’ *Philosophical Studies*, forthcoming.