MAY, MIGHT, AND IF

[1]

Table of Contents

[Introduction]

My main concern in this paper will be with the relation between semantic content and speech act force, with a focus on cases where a speech act aims at doing something other than communicating an item of information. I will try to bring out some connections and analogies between three controversies, each concerning a specific kind of modal construction. I will look first at David Lewis's language game of command and permission, second at the problem of epistemic modals, and third at debates about whether indicative conditionals express propositions.[2]

The classic discussions of speech acts distinguished the propositional content of a speech act from the illocutionary force with which that content is expressed. On the account of assertion that I have proposed, assertive force is explained in terms of the way the speaker intends to change the context, where context is understood as the common ground. Specifically, in the simplified and idealized assertion game, an assertion is a proposal to change the common ground by adding the content of the assertion to it. The change takes effect unless the assertion is rejected. This account of assertive force is characteristic of a wider range of speech acts than those that are properly called assertion, but assertive force, even in the broad sense that covers this range of cases, still differs from the force of some other speech acts that we will consider.

While the assertion game that I sketched uses a very abstract notion of force, it retains the traditional force/content distinction. I want to empasize the difference between a dynamic *pragmatics* and a dynamic *semantic* theory. The former says that content is determined as a function of context, and then a general force rule explains how the content changes the context. The latter streamlines the story, eliminating the middleman by building force into the semantic values of sentences. The projects I will discuss in this chapter all retain a distinction between content and force, in some form at least.

1. Commands and Permissions

I will start with what David Lewis called "a little language game" that he used to raise a problem about permission. Lewis's game was not presented as an analysis of any natural language expressions, but was the construction of an artificial game designed to throw light on the content/force difference. Traditionally, there are two ways to think about the role of modal notions, including the deontic modality involved in commands and permissions. On the one hand, one may think of a modal operator as determining a distinctive kind of proposition as a function of the proposition on which it operates (which, following the linguists, who are following medieval theorists, I will call the *prejacent*). But there is also a tradition of thinking of modality as determining the *mode* with which a proposition is expressed, and not being a part of, or a determinant of, the proposition itself. Lewis's game defines both a distinctive kind of deontic content, and a distinctive kind of imperative force, or mode, with which that content is expressed. First, the compositional semantics for the deontic sentences of the language that Lewis specifies determines propositions stating what one is obliged and permitted to do. Second, pragmatic rules are given for the use of the sentences in the playing of the language game. The pragmatic rules say how the things that the players might say will affect the state of play.

Specifically,[3] the language is a standard deontic propositional modal language, with sentence letters and two interdefinable modal operators, '!' and '¡' (which Lewis called 'fiat' and 'taif') for saying, respectively, what is required and what is permitted. In the standard Kripke-style semantics for such a language, a model consists of a set of possible worlds and a binary accessibility relation: the relation holds between worlds x and y if and only if world y is compatible with what is permissible in world x. Lewis's semantics is close to this familiar modal semantics, with world-time pairs playing the role of the possible worlds. Sentences of the form !ϕ and ¡ϕ, when used in possible world *w* at time *t*, say what is required or permitted at that time in that world. Their truth values, at ⟨*t*,*x*⟩, are determined by whether ϕ is true at some or all world-time pairs within a set of world-times that is defined in terms of ⟨*t*,*x*⟩ Lewis calls this set the 'sphere of permissibility' for ⟨*t*,*x*⟩.

There is little novelty in the language or in the truth-conditional semantics for it. The distinctive feature of the game is in the pragmatic rules, which are as follows: There are three players with distinctive roles: the master, the slave, and the kibitzer. The master controls the sphere of permissibility. When she issues a statement of the form !ϕ (at time *t* in world *x*), the sphere of permissibility adjusts to make the statement true at that time. More specifically, the master's command (at ⟨*t*,*x*⟩) makes it the case that the sphere of permissibility for time *t* in world *x* is the intersection of the sphere as it was before *t* with the proposition expressed by ϕ. The slave's job is to act in a way that ensures that the actual world remains within the sphere of permissibility. The kibitzer has no special powers or responsibilities, but can comment on the situation, perhaps reminding the slave what he must do, or the master what she has commanded or permitted.

The point of having a kibitzer in the game is to bring out the fact that the same sentence, with the same content, can be, in the mouth of one player, a command, and in the mouth of another an assertion. Both speech acts will have deontic content, but only one will have imperative force.[4]

The master can issue permissions as well as commands, and just as a command by the master makes it true that something is required, so a permission statement is supposed to make it true that something is permitted. In the case of a command, there is a straightforward answer to the question, what is the minimal adjustment to the prior sphere of permissibility necessary to make the content of the command true? So it is clear what the rule specifying the force of a command should be. But with a permission statement, there is no such answer. There is no unique way of "subtracting" a requirement in order to make a permission statement true. The task of finding a rule that says how the issuing of a permission statement by the master changes the context is the problem about permission that Lewis is referring to in the title of his paper. The paper does not offer a solution to the problem. It is argued that each of the proposed solutions that are considered either fails to be adequate, or is just a way of restating the problem.

I am not sure what Lewis would have counted as a solution to his problem (as contrasted with a restatement of it), since it is clear that any solution will require that new resources be added to the model, and any postulation of new resources sufficient to provide a determinate rule by which the spheres of permissibility evolve might be charged with being just a way of posing the problem. But a reformulation of the problem may still be helpful in sharpening it. Whether or not it counts as a solution, one can say something about the abstract structure of the resources that need to be added, and perhaps about the way the added resources might have application to other problems. Lewis's own work on the logic and semantics of counterfactual conditionals points the way, since there are parallels between the problem about permission and the problem of counterfactuals. The problem about permission might be put this way: suppose something impermissible was made permissible. What else would then be permissible? This looks a lot like the question: suppose a certain proposition that is in fact false were true. What else would then be true? The rough idea of Lewis's analysis of counterfactuals is that a counterfactual, (ϕ **£** $\rightarrow$ ψ) is true in possible world x if and only if ψ is true in all of the possible worlds in which ϕ is true *and which are otherwise minimally different from x*. A model for interpreting counterfactuals specifies a comparative "closeness" relation that determines what "minimal difference" comes to. One could add an analogous relation to the models for the commands and permissions game. There are alternative ways of making the structure precise, but one might suppose that the spheres of permissibility are nested in a succession of wider backup spheres, with some impermissible worlds "closer" to permissibility than others. When the master says '¡ϕ' at a given world-time, the prior sphere of permissibility expands to include those points at which ϕ is true, and which are in the "closest" sphere that includes some points at which ϕ is true. Nothing substantive could be said about the relevant notion

of closeness at this level of abstraction, but it seems reasonable to think that a permission game of the kind Lewis defined will require a relation with this structure, and if our models contained such a relation, it would yield a determinate rule for the evolution of the spheres of permissibility in response to permission statements by the master.[5]

It is this problem about permission that gets most of Lewis's attention, but there is also a small problem about commands (and permissions) that he does not consider, but that will be important for the analogy with epistemic modals that I will develop. The problem is this: The force rule for the master's imperative utterances is clear enough for simple commands – simple sentences of the form !ϕ. But Lewis's game gives a complete compositional semantics for sentences of the deontic language, including truth-functions of sentences that mix deontic and descriptive content. We need to consider what should happen when the master utters a sentence with mixed deontic and factual content. Suppose the master (Scrooge, in this example) says to the slave (Bob Cratchit), after having reluctantly given him permission to take Christmas off: "If you do take Christmas off, you must work an extra hour the next day." Or perhaps Scrooge uses a disjunctive sentence to perform his speech act: "Either you work on Christmas, or you must work an extra hour the day after." (If 'C' is 'Cratchit works on Christmas', and 'W' is 'Cratchit works an extra hour the next day' then Scrooge's sentence is (C ∨ !W))). What is the effect of the utterance of this disjunction? The general rule for the master's speech acts was supposed to be this: the sphere of permissibility adjusts in the minimal way required to make the master's sentence true, but it is not clear how to apply this rule to the disjunctive sentence with one imperative disjunct. If the sphere contracts to require that the Cratchit is (categorically) required to work the extra hour, this will ensure that Scrooge's utterance is true, but this is too strong, since the statement might be true without this change. On the other hand, if the sphere contracts only enough to require that Cratchit is obliged to make the disjunction (C ∨ W) is true, this will not suffice to ensure that Scrooge's sentence, interpreted in a straightforward way, is true. Suppose Cratchit fulfills his obligation by taking Christmas off, and then working the extra hour the next day. Scrooge's statement will then be false since the first disjunct will be false, and because there is no categorical requirement, the second will be false as well.

The literal-minded Cratchit might be tempted to reason this way: "Either the command disjunct is true, in which case I have to work the extra hour whether or not I take Christmas off, or it is false, in which case I have no extra obligations, no matter what I do. In the former case, I prefer to take Christmas off rather than work Christmas, and also the extra hour the next day. In the latter case, I also prefer to take Christmas off (since in this case, I still won't have to work any extra hours). So in either case, that is what I should do. Of course I recognize that since I plan to take Christmas off, it must be that *if* the Master's command is true, it is the second disjunct that is true. but perhaps the master's statement is false, since it is at least *partly* a prediction, and therefore not automatically true. Furthermore, since I know that the master has not issued any new categorical command, I have reason

to think that the second disjunct is false, whatever I do, and so (given that I am planning to falsify the first disjunct) it seems that the master's statement is just mistaken." Has the clever Cratchit found a loophole that will allow him to have his Christmas off, with no extra hours the next day? Scrooge might protest, "What I *meant* to say was that you must either work Christmas, or work an extra hour the next day." "But," Cratchit replies, "that is not what you said."[6]

Of course Lewis's game is just a made-up exercise, and we can adjust either the compositional semantics or the force rules in any way we like. One might simply stipulate that an utterance from the master counts as a command (or a permission) only if it is an unembedded deontic sentence. On this proposal, a sentence with an embedded deontic clause will always be treated as an assertion, even in the mouth of the master. On this stipulation, the reasoning that we put into the mouth of Cratchit is perfectly correct. To avoid this consequence, one might stay with this semantics, and this policy, for the official language, but allow commands to be made with *surface* sentences with imperative parts so long as they can be reinterpreted and formalized as categorical imperatives (for example taking the real form of "either C or ought-W" to be '!(C $\vee$ W)'). But these moves would rob the game of much of its interest, which is the promise that it can combine the advantages of a compositional semantics with a general account of speech acts that are used to do something different from what straightforward assertions do, which is to state what the world is like. The hope is that a game with this feature will throw some light on the way deontic and other modal expressions work in natural language. For this reason, the way we adjust the rules of the artificial game to solve the problem should be guided by the related phenomena in natural language, and it seems intuitively that complex sentences with embedded deontic clauses can be used (by a speaker with the right status) to change what is required. The problem for the game, and the adjustment to solve it, should help to explain why the natural language constructions that are the analogues of the operators in our simple game have some of the features they have.

The solution I will suggest modifies, not the force rule, but the compositional semantics, and it will make use of the notion of a subordinate context that is independently motivated by its use in the explanation of phenomena concerning presupposition and anaphora. The simplest case of a subordinate or derived context is a temporary context created by a supposition. In the case of an indicative suppositions, the derived context is represented by the subset of the set of basic context set in which the proposition supposed is true. Subordinate contexts or this kind are wholly defined in terms of basic contexts, so in a sense they are features of the basic context – the common ground. One central role of subordinate contexts is to help to explain the dynamic process of presupposition accommodation, and to bring into focus the information provided by the basic context that is available for determining the propositions expressed or denoted by embedded clauses. But there is no reason why a language could not have modifiers or operators that are interpreted by semantic rules that appeal to relevant subordinate contexts, when the operators or modifiers occur in

embedded clauses. If we interpret the deontic operators in Lewis's language this way, we can explain why some complex sentences with imperative parts can have the structure that they appear to have (a disjunction with one imperative disjunct, or a conditional with an imperative consequent), and also be governed by the general force rule when they are used by the master. This is the rule that says that the sphere of permissibility adjusts in the minimal way required to ensure that the sentence as a whole is true. So suppose we add to the semantics for Lewis's game an extra parameter of the interpretation, a *context*, represented by a set of world-time pairs. The semantic rule for the imperative, in the unmodified theory, was this:

$[!\phi]\langle t,x\rangle = 1$ iff $[\phi]\langle y,t\rangle = 1$ for all $y \in S\langle t,x\rangle$

(where $S\langle t,x\rangle$ is the sphere of permissibility at $\langle t,x\rangle$)

The modified rule will be this:

$$[!\phi]\langle t,x, C\rangle = 1 \text{ iff } [\phi]\langle y,t, C\rangle = 1 \text{ for all } y \in S\langle t,x\rangle \cap C.$$

(where C is the new parameter, a set of time-world pairs)

The default or initial context (for the interpretation of unembedded sentences) will be the set of all the relevant time-world pairs, "all of those such that *t* is a time during the game, and *w* is accessible (at the actual world) at the time at which the game begins."[7] So the modified rule will yield the same result as the original rule for unembedded commands. But the context parameter may be shifted by a compositional rule for another operator. With this modification to the semantics, the force rule can be general: the master's utterance of the sentence with the embedded deontic clause will adjust the sphere of permissibility to ensure the truth of the complete sentence.

It is not that every sentence with an embedded deontic clause will be suitable for issuing commands. There is no way, for example, that a belief attribution to someone other than the master could be used, by the master, to change the sphere of permissibility. If the master said to the slave "The kibitzer believes that you must work an extra hour on the day after Christmas," this will be interpreted as an ordinary assertion, since no change in sphere of permissibility would suffice to make it true.

Though the suggested modification of the semantics is motivated by the role that we want imperative and permission statements to have in the overall game, the semantics remains independent of the force rule, and the modified semantics is still static in the sense that the semantic value of any sentence in a given context, and at a given time, is a function of the semantic values of the parts at that same context and time.

We haven't said exactly how the parameter shifts in a disjunctive or conditional context – that will depend on the semantics and pragmatics of disjunction and the conditional. Conditionals I will discuss later. Disjunction is complicated, with the order of the disjuncts sometimes mattering, and sometimes not. (It seems to matter in this case). The parameter shift might in some cases be optional, allowing for scope distinctions, though some possible interpretations that make sense for the kibitzer's assertions might not make sense for the master's commands. If the shift is optional, then even with the modified rule, the disjunctive sentence, in the Scrooge example, might have the interpretation that Cratchit took it to have. The kibitzer, for

example, might remember that either the master commanded the slave to work an extra hour, or else the slave decided to take Christmas off, but can't remember which. But the master could not use the sentence, interpreted this way, to issue a command.

There is more to say about the ways that Lewis's game might be refined and extended, but I will leave that for another time. To conclude this section I will point to some features of the game that are relevant to our general concern with the relation between content and force.

To the question, "Should the distinctive character of a command or the issuing of permission be represented with a modal operator, affecting the content, or by a force operator, determining the way that a given content affects the context?" Lewis answer was "both." Deontic content is explained independently of force, so we can account for propositions about requirements and permissions that are embedded in complex statements, and explain how propositions about requirements and permissions can be the content of speech acts of different kinds. But the theory also gives a precise account of a distinctive kind of imperative force, and an explanation of the relation between the distinctive kind of content and the distinctive kind of force. The semantic account of the content is *static* in the following sense: The compositional rule for a deontic sentence, !ϕ or ¡ϕ, specifies the proposition expressed at a given point (world-time pair) as a function of the proposition expressed by the prejacent ϕ at that same point.

The pragmatic force rule is like the rule for assertion in that it explains force in term of the way that the speech act changes the context, as a function of the content of the speech act. But commands and permissions issued by the master contrast with assertions (including assertions of deontic propositions) in this way: they change the context by changing a parameter of interpretation (the sphere of permissibility) so as to make the proposition true. The result is that in these cases the determination of the content of the speech act is *prospective* in the following sense: the content of the deontic sentence is interpreted relative the *posterior* context – the context as it is after the force rule is applied. Because of this feature of imperative force, it contrasts with assertive force, where the content of the assertion is determined relative to the *prior* context. The prior context is not the context as it was before the speech act took place, but the context as it is after the manifest fact that the speech event has occurred is recognized. But the context change, in the case of the imperative speech acts of Lewis's game, is not a case of presupposition accommodation. In these cases, it is the application of the force rule itself that brings about the change in the parameter of the context relative to which the content of the speech act is determined.

The form of this dynamic rule (context is changed as a function of a proposition that is determined as a function of the context that results from the change) introduces an obvious threat of circularity. The strategy for avoiding circularity is to specify that the context is to be changed in the minimal way required to make the proposition true. In the case of commands, it is clear what the unique minimal change is, but in the case of

permission, additional structure must be added to the models in order to get a determinate force rule that avoids circularity.

So let me sum up the lessons of Lewis's little language game, lessons that will carry over to the other modal expressions I will consider: First, the semantics is static, and is explained independently of force. Second, the way imperative speech acts change the context is prospective. Third, the prospective character of the force rule requires (in the case of permission) that structure beyond that required for the static semantics be introduced in order to give a determinate account of the dynamics of the game.

2. Epistemic modals

The traditional view about epistemic 'might' and 'must' is that they quantify over the possibilities that are compatible with the knowledge or potential knowledge of some contextually specified individual or group. When Alice says to Bert that it might rain, she says that she doesn't know that it won't rain, or perhaps that neither she nor Bert knows it, or that she is, or they are, not in a position to know it or to come to know it. The puzzle is that no way of pinning the relevant knowledge state down seems to be able to explain both why we are in a position to make the epistemic 'might' claims we seem to be in a position to make, and also why it is often reasonable to disagree with 'might' claims made by others. If Alice's statement were about just her own state of knowledge, then how could Bert disagree with her (as it seems he could)? On the other hand, if Alice's claim is about what information is available or potentially available to some larger group, then Alice will not be in a position to know what might be true in cases where it seems, intuitively that she is. Different versions of the traditional view – what von Fintel and Gillies have called "the canon"[8] – try to find a way to pin down the constraints on the relevant informational state that can explain the data.

Seth Yalcin's account of epistemic modals[9] began with a striking observation that raised a new puzzle. First, he observed that, as one would expect on the canonical view of epistemic modals, there are versions of Moore's paradox for epistemic modals that parallel Moore's paradoxes for explicit knowledge claims. Just as one cannot coherently say, "it will rain, but I don't know that it will," so one cannot coherently say "it will rain, but it might not." In the case of explicit knowledge claims, the paradoxical conjunctions are shown not to be *semantic* contradictions by the fact that they can be coherently ***supposed*** even if they cannot be asserted. There is no paradox or incoherence is saying, "Suppose it is will rain, but that I don't know that it will," or "if it will rain and I don't know it, then I will get wet." But with epistemic modals, the incoherence persists in the supposition context: "Suppose that it is raining, but that it might not be raining" is just as bad as the simple conjunctive assertion. So the conflict between 'is' and 'might not' is deeper than in the familiar cases of Moore's paradox. [10] The new puzzle shows that 'ϕ, but it might be that not-ϕ' cannot simply mean, 'ϕ, but not-ϕ is compatible with X's knowledge', for any choice of X, since whatever the choice of X, it should be perfectly coherent to suppose that ϕ is true, but that X does not know it. But despite this fact, 'ϕ, but it might be that not-ϕ' still cannot be a straightforward contradiction, since that would imply that the claim that it might rain entails that it will rain, which is obviously wrong.

Yalcin offers a detailed semantic and pragmatic account of epistemic modals that is designed to solve both the traditional problem about bare epistemic modals, and his new problem about epistemic modals embedded in contexts of supposition, and in propositional attitude ascriptions. I will exploit some of the insights of his account, but will spell out my take on the problem in my own way, which will emphasize the analogy with Lewis's game of commands and permissions. I follow both Yalcin and Lewis in

separating the task of giving a compositional semantics for the relevant modal expressions from the task of explaining the pragmatic role of the resulting semantic value. And I will follow them in doing the semantics in the familiar truth-conditional framework. Despite the truth-conditional form of the semantics, Yalcin describes his account of epistemic modals as *expressivist*, and the basis for that description is that the speech act that epistemic modals sentences are used to perform (according to the account) contrasts with the speech act of assertion, which aims to describe the world as being a certain way. Lewis's account of the speech acts of the master in his language game might be called expressivist for the same reason.

As we saw, what is distinctive about the force of the master's speech acts in Lewis's game is that the content of the command or permission sentence is determined by the *posterior* context, rather than the *prior* context, as is the case with assertions. My proposal is to make the same move in explaining the distinctive force of a statement made with an epistemic modal. In the master/slave game, the relevant feature of the context is the sphere of permissibility. In the case of epistemic modals, the changing contextual feature will be the context set itself.

As with Lewis's master/slave game, we start with the static compositional semantics, and as with the semantics for Lewis's game, the semantics will be close to an orthodox Kripke semantics (at least if we ignore certain complications that we will consider later). Our framework already provides us with a binary relation that is determined by a common-ground context. The context, in this sense, is a set whose elements are possible worlds centered on a time and a sequence of individuals, the participants in the conversation. The conversation is taking place in each of the possible worlds in the context set, at the time, and involving the parties that define the center. So for any world *x* in the domain of the common ground, there is a set of worlds that are compatible with what is common ground in *x*. So this gives us a binary accessibility relation on s-worlds, R, defined as follows: *x*R*y* if and only if *y* is compatible with what is common ground in *x*. If we restrict our domain, for the moment, to possible worlds compatible with the common ground of a given conversation, and assume that our context is nondefective in the sense that the parties to the conversation are all making the same presuppositions, then the accessibility relation will be an equivalence relation, and so the logic of the 'might's and 'must's interpreted by it will have a simple S5 structure. (We will come back later to the question what happens when we interpret an epistemic 'might' or a 'must' in a possible world outside of the context set). That is all there is to the semantics. As with Lewis's game, the innovation comes when we specify the force rule, and we will start, as Lewis's game does, just with a rule for unembedded modal claims. The proposal is to make the rule *prospective* in exactly the sense in which the commands and permissions issued by the master in Lewis's game are prospective. In saying might-ϕ, one is not *asserting* that ϕ is possible, relative to the *prior* context. Rather, one is proposing to adjust the context (if required) to bring it about that what the sentence says, relative to the *posterior* context – the context as adjusted – is true.

In Lewis's game, one player – the master – had complete authority over the relevant parameter, and so that parameter automatically adjusts in response to her speech acts. But in the assertion game, no one player controls the context set; it is subject to negotiation, and this will lead to complications not found in Lewis's simple game. The slave does not have the authority to reject the master's commands or permissions, but epistemic 'might' statements, like ordinary assertions, can be rejected. Since the context set models what is commonly accepted, if one party refuses to rule out a possibility, it must remain in the context set. But parties to a conversation can still disagree about whether they are in a position to rule out a possibility. I will defer the question what happens in cases of disagreement (discussed in a later chapter of the book from which this paper is drawn). Let's consider for now a standard and unremarkable case where the speech act changes the context, but does not result in disagreement. Alice says "Noam might be in his office," and Bert responds, "No, I just saw him leaving for lunch." Alice's statement was true, relative to the *prior* context, but we are taking it to be a proposal about the posterior context, a proposal that possibilities compatible with Noam being in his office remain in the context set. Bert rejects the proposal, since he has information sufficient to rule these possibilities out.

If it was already compatible with the prior context that Noam is in his office, why did Alice need to say what she said? In the assertion game, we assume that it is inappropriate to assert what is already accepted, since the assertion would have no effect on the context. This question points to two related ways in which we need to refine the simple picture of the common ground, ways that are independently motivated. First, as Kripke's discussion of presupposition brought out, there is a difference between information that is general background knowledge and information that is presupposed in an active context. Assertions may be appropriate as reminders, even if they are not really news.[11] Such assertions serve to bring an item of background knowledge into the active context – to make it salient. Second, both philosophers and linguists have recently emphasized the importance of recognizing the questions that are at issue in a given context[12], and a representation of context needs, not just a space of those possibilities that are compatible with what is presupposed, but also a partition of the possibility space, with the questions at issue being those that distinguish between the cells of the partition. If we add this kind of structure to our representation of the common ground, then there will be two ways in which a 'might' statement can change the context. In some cases, it will expand the space of relevant possibilities in the minimal way required to make the 'might' statement true, relative to that revised context. But in other cases, the job of the 'might' statement will be to refine the "modal resolution" (to use Yalcin's term) – to add a new distinction between the possibilities. The following example is a clear case of the first kind of change: Bert says, "The butler didn't do it, so it must have been the gardener", presupposing that the guilty party was one of the two. Alice replies, "Wait, it might have been the chauffeur – we forgot about him." To illustrate the second kind of context change. Yalcin uses this example: The speaker says, "it might be raining in

Topeka." It was not previously presupposed that it was *not* raining in Topeka – the issue was not on the table. The statement neither made nor retracted a claim, but served only to raise the question.

Because of the prospective character of the force rule, the interpretation of those 'might' statements that are not compatible with the prior context faces a problem that parallels Lewis's problem about permission. Without additional structure beyond that required by the static semantics, we don't have determinate context change rule for the epistemic 'might'. The kind of structure we need is the same as that required for Lewis's game: nested spheres of possibility around the basic context set. This is exactly the kind of modal structure that semanticists such as Kratzer and von Fintel have used in their accounts of modality and conditionals.

The problem about embedded commands and permissions that I raised for Lewis's game also has a parallel in the epistemic case, and the solution is the same. If Alice were to say, "if it rained, the game must have been called off," or "even if the Yankees lost, they might still win the division," how is she proposing to adjust the context? Unless we modify the semantics, no adjustment to the basic context will be both minimal and sufficient to make the statement true. But if we interpret the 'must' and 'might' relative to a derived context, the complex sentence can be a proposal of the right kind: one put forward as true relative to the posterior (basic) context.

This modification gives us a solution to Yalcin's puzzle about 'might' in the scope of a supposition, a solution that is close to (and modeled on) the one that he proposed. Yalcin's theory, formulated in Lewis's Context/Index framework, proposed that an information state (represented by a set of possible worlds) be added to the index, and this coordinate will be "shifted" for the interpretation of embedded clauses. If we formulate my suggestion in the this framework, it would add a set of (multiply-centered) worlds to the index. I don't agree with Yalcin's original proposal about the way this coordinate of the index shifts in the scope of propositional attitude ascriptions, for reasons that his later work has brought to light, but I agree with the basic idea, and I think the prospective form of the force rule helps to motivate it.[13]

The move that Lewis made in defining his language game (with the adjustments I proposed) allowed for the combination of a smooth compositional semantics with a force rule that made sense of speech acts that do something different from conveying information. On this kind of account, nonassertive speech acts and assertions have in common that a proposition determined by the semantics is put forward as true. In both cases, the end result of a successful and uncontested speech act is a context in which the content of the speech act is accepted – true in all possible worlds in the context set. The contrast between assertions and the other speech acts is in the means by which this is accomplished.

Neither Lewis nor Yalcin discuss the case of deontic or epistemic modals in the scope of quantifiers, but one advantage of a theory that uses a truth-conditional semantics, even for determining the contents of nonassertive speech acts is that it can give a smooth semantics for such cases. There could be a version of Lewis's game with many slaves, and in such a game,

the master could issue general commands and permissions such as this: "each of you may take a day off on your birthday." With epistemic modals, we have examples such as this: "Any one of the candidates might be selected." (Though with both the Lewis game, and the parallel proposal about epistemic modals, not every quantifier would be suitable for issuing permissions, or for altering the common ground. If the master were to say, "some of you may take a day off on your birthday," this would have to be interpreted as an assertion. If Alice were to say "several of the candidates might be selected," this will require (or at least favor) a wide scope reading for the 'might' (it might be that several of the candidates are selected), or else a nonepistemic interpretation of the 'might'.)

In giving the semantics for the epistemic 'might' and 'must', we restricted the specification of the accessibility relation to worlds that are compatible with the given context. But what about possible worlds outside of the context set? For example, Alice says that it might rain tomorrow in a context in which it is presupposed that the weather forecaster predicted rain. Now consider a possible world *w* in which the weather forecaster did not predict rain (perhaps no prediction was made) and in which it in fact did not rain on the relevant day. Is Alice's 'might' statement true or false in that possible world? Since presupposition and common ground are not factive notions, it could be that Alice's presuppositions in world *w* are the same as they are in the actual world. In fact, *w* might be the actual world. So we know what the context set is, in world *w*, but what is the set that is R related to world *w*, where R is the relation that is relevant to determining the truth-value of the epistemic modal sentences in that world? One might be tempted to say that the relevant R is just the relation that holds between a world *w* and the worlds that are compatible with what is presupposed in *w*, and this would be the right thing to say if the modal statement were simply the proposition that something is compatible with what is being presupposed. But there is a subtle difference between a proposition that states something *about* the context and a proposition that expresses something that is context-dependent. ϕ is compatible with the common ground if and only if we are presupposing that it might be true, and that is enough to tell us whether the 'might' statement is true or false in the worlds compatible with the context, but it leaves open the question, whether the proposition expressed in that context is true or false (or neither) in possible worlds outside the context. So our theory, thus far, remains silent on the question whether the 'might's and 'must's are true or false (or neither) in worlds outside of the context set. Should this gap be closed, and if so how?

To get a more concrete sense for this kind of question, let's look at some examples. In many cases, there will be a clear and unproblematic fact of the matter about whether some statement is true or false, even if the statement is made in a context in which a relevant presupposition is in fact false. Suppose I try to call Sam Smith, but dial the wrong number and am in fact talking to someone else. "Is this Sam?" I ask, and because my interlocutor happens by coincidence to have the same first name, he answers "yes". Presupposing that I am talking to Sam Smith, I say, "The department voted to offer you the job." It is clear enough what I have said: I referred, with my

use of 'you' to my actual interlocutor, and inadvertently said something false about him. But in other cases, the semantics may not provide an answer, or it may be controversial what the answer should be. Suppose my interlocutor and I are presupposing of Jones that he is the unique man drinking a martini and I say "The man drinking a Martini is a philosopher."[14] Jones is in fact a philosopher, but our presupposition that he is drinking a martini is false – it is Perrier in his martini glass. Russell, Strawson and Donnellan will all agree about the way the assertion affects the context, since they all agree about the truth-value of the statement with the definite description in possible worlds compatible with what is being presupposed in the context. But their analyses of definite descriptions will disagree about the truth-value of the statement in the *actual* world, which is outside of the context set: Assuming that no unique man in the room is actually drinking a martini, Russell will say the statement is false, Strawson will say that it is neither true or false, and Donnellan (assuming the description is being used referentially) will say that it is true.

On the dynamic *pragmatic* story, it is a minimal requirement on appropriate speech that one use a sentence that has a truth-value *at least* for all the possible worlds compatible with the common-ground context. That is, it is required that the function from possible worlds to truth-values determined by the semantics be a total function, relative to that domain. On the dynamic *semantic* story, where the semantic values are context-change potentials, truth-values are determined *only* for possible world in the context set. On the pragmatic theory that separates the determination of content from the force of the speech act, the proposition expressed *may* extend beyond the possibilities compatible with what is being presupposed. On this kind of account, some speech acts are more robust: what is asserted may be detached and added to one's stock of beliefs. In other case, what is said may be more fragile, succeeding in distinguishing between the possibilities compatible with a local context, but not extending much or at all beyond it. The silence of the account we are considering about the relevant information state for interpreting epistemic modals in possible worlds outside the context set implies that sentences with bare epistemic modals will, or at least may, be cases for which the proposition expressed is partial, and so the statement gets no truth-value in the actual world. But in some cases, there may be a natural extension of the relevant parameter to worlds outside the context set. Consider this extended example. It is the last day of the baseball season, and the Yankees are one game ahead of the Red Sox. Each team has one game to play, though not with each other. If the Yankees win, or if the Red Sox lose, the Yankees win the division. But if the Yankees lose *and* the Red Sox win, they will then be tied, so there will be a playoff game to determine the division champion. All of this is common ground in the context. Alice says: "The Yankees lost, so the Red Sox still might win the division." This is true in all of the possible worlds compatible with the (posterior) context. But, let us suppose, Alice was wrong about the Yankees – they in fact won. Alice's 'might' statement was appropriate, and succeeded in changing the context in a determinate way, but was it true *in the actual world*? *We* know that the Red Sox won't be the division winner,

but we have information not available to Alice and her interlocutors. It could have been that one of the interlocutors knew that Alice was mistaken, in which case he would reject both her assertion about the Yankees, and her proposal that possibilities in which the Red Sox win the division remain in the context set. But no one in the conversation had this information. In this case, was Alice in fact wrong – not just about the Yankees game, but about her epistemic 'might' statement about the Red Sox? Now consider the reverse situation: Suppose Alice said, "The Yankees won, so the Red Sox must be out of the running for the division title." Alice got it wrong in this version of the example too: the Yankees in fact lost. But alas, the Red Sox also lost (though none of the relevant parties know this), so the Red Sox *are* in fact out of the running. Is Alice's statement that they *must* be out of the running true in the actual world? Her reasoning was based on a false premise, but is this enough to make it false? I don't think these questions have obvious answers. There is some inclination to fill the vacuum left by our theory's silence by shifting to *our* context – the context of the theorist giving the example, bringing in the stipulated facts about the story that are used to raise the problem. There is also some inclination to shift to a nonepistemic modality to give an answer. (If the Yankees lost, and the Red Sox game is still undecided, then even if we stipulate that they *will* lose, we can say that, as of now, they still *might* win. But if we say this, we are using the 'might' in a nonepistemic sense.) Further elaboration of the example, or of the theory, might give reason to assign a truth-value, outside of the context set in such cases, but for some cases, the semantics might resolutely retain its silence.[15]

The semantics will give a definite answer, even outside the context set, in the case where the prejacent of the 'might' statement is true in a given world, or when the prejacent of a 'must' statement is false. Whatever anyone believes, knows or is presupposing, if the Red Sox actually did win the division in the end, then the statement that they might win was actually true, relative to any context in which that statement was made, or denied. That follows from the reflexivity condition on the relevant accessibility relation, which it seems intuitively clear should hold. And this condition will hold for deontic 'must' and 'may' as well as for the epistemic modals. Lewis's little language game used an artificial language, with '!' and '¡' for the command and permission operators, but if Lewis's commands and permissions were made with 'must' and 'may', with their normal senses, then a disobedient slave will be a problem, not just for the master, but for the semantics. The master can make it the case, just by her speech act, that the slave is obliged to stay out of her wine cellar, but she cannot make it the case, by her speech act, that the slave does what he is obliged to do. If either the master or the kibitzer says to the slave, "you must stay out of the wine cellar," she or he must presuppose that the slave will do what he must do. If one wants to allow for the possibility of the disobedient slave, one must put the command or statement differently – for example, with 'should' or 'ought' rather than 'must.'[16]

3. Indicative conditionals

Debates about whether indicative conditional statements express propositions predate the recent flurry of discussion about epistemic modals, but the parallels are clear. I will argue that we can reconcile the propositional and non-propositional approaches, and my attempt to do this will replay the same themes seen in our discussion of deontic and epistemic modals. On the one side of the old debate about conditionals is the hypothesis that a conditional statement, whether indicative or subjunctive, expresses a proposition that is a function of the propositions expressed by its constituents, the antecedent and the consequent. The function may, and is usually assumed to be, context-dependent. On the other side is the hypothesis that the indicative conditional is used to make a distinctive kind of speech act involving just two propositions – those expressed by antecedent and consequent. A conditional assertion is not the unqualified assertion of a conditional proposition, but a qualified assertion of the consequent, with the antecedent expressing the qualification. On this hypothesis, the conditional sentence does not purport to state a fact, but expresses an epistemic attitude, or makes a qualified commitment. Dorothy Edgington has been the most prominent proponent of the conditional assertion account. Building on earlier work by Ernest Adams, she has given both challenging arguments against the proposition analysis, and a constructive development of the view. Allan Gibbard has also argued for what Jonathan Bennett labeled the NTV (no truth-value) analysis of indicative conditionals, which Bennett also endorses. Gibbard and Bennett accept a propositional analysis of ***subjunctive*** conditionals, arguing that the two kinds of conditionals should receive separate analyses, while Edgington opts for a unified account according to which even counterfactuals should be understood as sentences that express certain probabilistic relations but do not express propositions. On the propositionalist side, some philosophers and linguists have defended the material conditional analysis, for indicative conditionals, attempting to explain away apparent counterexamples as cases that are true, but conversationally inappropriate. Grice adopted this strategy in his William James lectures. David Lewis, Frank Jackson, and Barbara Abbott have also argued for this analysis. Others have argued for a unified truth-conditional analysis for both kinds of conditionals. I have been on both sides of the debate between the propositionalists and the proponents of a conditional assertion account, first arguing, in a paper published in 1975[17], that the same abstract possible-worlds semantics developed for the interpretation of counterfactuals should also be used for the interpretation of indicative conditionals, with the semantic differences between the two kinds of conditionals explained by different constraints on the contextually determined parameter of the interpretation. I still think this is the right approach, though I later acknowledged, in response to arguments by Allan Gibbard, that if indicative conditionals are to play the roles that they seem to play, they "must be too closely tied to the epistemic states of the agents who utter them to express propositions which could be separated from the contexts in which they are accepted."[18] In a recent paper, I argued that one could reconcile the conditional assertion analyses with a propositional

account, seeing the former as equivalent to a special case of the latter. The general strategy parallels the one we have seen in the language game of commands and permissions, and in the semantic/pragmatic account of epistemic modals. In all of these cases, the aim is to combine the advantages of the compositional semantics that comes with a propositional approach with an acknowledgement that the expression in question is being used to do something other than to communicate an item of information.

To spell out how this idea applies to indicative conditionals, I will first review the semantic/pragmatic analysis of indicative conditionals that I have defended, and the spell out the sense in which one version of this account coincides with a conditional assertion analysis. The proposal will be that indicative conditional statements are *prospective* in exactly the way that permissions, commands, and epistemic 'might' and 'must' are prospective. After sketching the idea as it applies to simple conditional statements, I will look at indicative conditionals within the scope of quantifiers. As with the other modal expressions, quantified cases illustrate the advantages of the propositional form of analysis, but they also bring out problems raised by the interaction of the semantics with the pragmatic constraints. The discussion of conditionals will continue in the next chapter, where I will consider the relation between the fact-stating role of conditionals and their role in expressing epistemic attitudes.

On the propositional account of conditionals that I have defended,[19] the semantics interprets the conditional in terms of a context-dependent selection function taking a possible world *w* and a proposition ϕ to a possible world, $f(w, \phi)$ – intuitively, a world in which ϕ is true, but which otherwise differs minimally from *w*. Nothing substantive is said in the abstract theory about the criteria for minimal difference, but formal constraints are put on the selection function to ensure that it orders the possible worlds, with the base world, *w* first in the order (since any world is minimally different from itself). The semantic rule is a follows:

$(\phi > \psi)$ is true in possible world *w* if and only if ψ is true in $f(\phi,w)$.

This semantic rule gives the truth-conditions for any conditional in a given possible world, and it also determines the *subordinate context* for the conditional supposition. If C is the context set (the set of possibilities compatible with the common ground), then the subordinate context for the supposition ϕ is $\{f(w,\phi): w \in C\}$.

For indicative conditionals, the pragmatics adds a contextual constraint on the selection function. Where C is the context set, say that a selection function is *admissible* if it meets the following condition:

If $w \in C$, then $f(w, \phi) \in C$.[20]

The effect of the constraint is to ensure that all of the presuppositions of the basic context are preserved in the subordinate context.

That is the truth-conditional semantics. We can also formulate, very simply, a conditional assertion analysis in our pragmatic framework, as follows: There is a speech act of *supposition*, which like the speech act of assertion, adds the content of the speech act to the common ground, but in this case it is added temporarily, and will no commitment to the truth of the supposition. The consequent of the conditional is then asserted in this

temporary subordinate context. At the end, the possibilities that were temporarily removed are added back.[21]

Though one of these stories is propositional and the other is not, they have much in common. Each can explain why an indicative conditional assertion is very similar, in its effect, to the assertion of a *material* conditional, but each can also avoid some of the problems with a simple material conditional analysis. For example, each can explain why the rejection or denial of a conditional is very different from the rejection or denial of a material conditional. The similarity between the two analyses can be made precise by defining a version of the propositional analysis that is essentially a terminological variant of the conditional assertion view.[22] First, we note that any contextual parameter relative to which truth-conditions are specified (a modal base, a domain of discourse, a selection function, the referent of a demonstrative) may be underdetermined by the actual context. In such cases, a proposition may be partial – a truth-value determined for some, but not all possible worlds. Suppose we adopt a version of our propositional semantics for indicative conditionals that is maximally cautious, determining a truth-value for a conditional only when the rules we have specified suffice to determine one, as a function of the antecedent, the consequent, and the common ground. This is accomplished by saying that all *admissible* selection functions are on a par. A conditional is true (relative to a given CG-context) if true for all admissible selection functions, false if false for all, and neither true nor false otherwise.[23] The upshot will be that a conditional sentence will be false in possible worlds compatible with the context if and only if that possibility would be excluded by the conditional assertion, according to the conditional assertion analysis. So one can see this version of the propositional theory as an implementation of the conditional assertion account. But this will work only if we assume that the interpretation of the conditional sentence is *prospective*. Here is the argument that this assumption is necessary: suppose I *believe* that learning ϕ would be sufficient reason (given my other beliefs) for accepting ψ, but that this is not common ground, since $(\phi\&\sim\psi)$ is compatible with the prior context. A conditional assertion would clearly be appropriate in this situation, since in a context in which ϕ is supposed, I am in a position to assert ψ. But on the maximally cautious version of our truth-conditional semantics, if we interpret the conditional (if ϕ, then ψ), relative to the *prior* context, the conditional will be neither true nor false in the possible worlds in which ϕ is false. Since I am not in a position to rule out these possibilities, the conditional, interpreted this way, would not be assertible. But if we interpret the conditional relative to the *posterior* context, where that context is the result of the minimal adjustment necessary to make the conditional true, relative to the adjusted context, then the conditional statement will be appropriate, and we get the equivalence between our two accounts. Just as with the epistemic 'may' and 'must' statements, the effect is to adjust at once a parameter of the interpretation and the set of possibilities compatible with the common ground. In the epistemic modal case, the changing common ground *is* the parameter of the interpretation; in the case of conditionals, the parameter is closely constrained by the

changing common ground. The general account explains both how the expressive speech acts are like assertions, and how they are different. In all cases, the end result of a successful speech act will be a context in which the proposition expressed is true in all of the possible worlds compatible with what is then the common ground. In the non-assertion cases, this success is achieved, in part, by changing the determinants of the truth-conditions for the proposition.

If the upshot is the equivalence of the conditional assertion analysis and (one version of) the propositional account, why is it important to go with the truth-conditional semantics? Part of the answer is the familiar one – that it provides a semantics for sentences in which conditionals are embedded in other constructions – but of course this is an advantage only if the compositional semantics gets the right result. Lewis appealed to compositionality in defending the presumption that indicative conditionals express propositions . "We think we know how the truth conditions for compound sentences of various kinds are determined by the truth conditions of constituent subsentences, but this knowledge would be useless if any of those subsentences lacked truth conditions."[24] But the particular analysis of indicative conditionals that Lewis defended was the material conditional analysis, and it is precisely with the embedded conditionals that this analysis gets things wrong. As we have seen, simple assertions of indicative conditionals have essentially the same effect as the assertion of the corresponding material conditional, on both the conditional assertion analysis, and on our minimal truth-conditional hypothesis, but the most dramatic divergence between ordinary indicative conditionals and material conditionals occur with denials, rather than assertions, or with conditionals embedded under negation. The truth-conditional account I am defending does better with embedding under negation, and I will argue that it also does better with conditionals in quantified sentences.

A second advantage of a truth-conditional semantics for indicative conditionals is that it allows for continuity between cases where conditionals seem to express an epistemic attitude and cases where their aim seems to be to communicate information about the world, information that is independent of the epistemic situation of the participants in a conversation in which the information is exchanged. It may be a matter of debate and negotiation, not only what the facts are, but what there is a fact of the matter about, and it is good to have a framework that does not require that such questions be settled in advance. I have more about this issue in another context, but for now let's look at the compositional considerations – particularly at the interaction of indicative conditionals with quantifiers.

Jim Higginbotham argued that quantified conditionals present a prima facie case of non-compositionality.[25] He begins by comparing two examples, with different quantifers, each with the a conditional open clause in the scope of the quantifier:

(1) Everyone will succeed if he works hard

(2) No one will succeed if he goofs off

He claims that the first can be interpreted as a conditional in the scope of a quantifier, but that the second cannot. Higginbotham is not explicit about how he is interpreting the conditional at this point in the discussion, but what he says about (2) suggests that his noncompositionality claim is based on a strict conditional interpretation. Suppose we interpret (2) as saying that for no x (in the relevant domain) is it true that if x goofs off, x succeeds. His claim is that it would then say that goofing off is in no case a sufficient condition for success, while the actual meaning is that in no case is goofing off compatible with success. The point is all the more obvious if we interpret the conditional as a material conditional, for then (2) would assert that everyone will goof off, and that no one will succeed, which is obviously not what it says. Higginbotham then suggests that compositionality can be restored if we give the semantic analysis that I proposed, and that I have sketched here, though he does not consider the pragmatic constraint. On this analysis, one could paraphrase (2) this way:

(2') The following is true of no x (in the relevant domain): in the world minimally different from ours (in relevant respects) in which x goofs off, x succeeds.

This seems to get things right.

Higginbotham considers an alternative analysis that takes 'if' clauses to be restrictors on the quantifier, rather than sentential connectives. On this analysis, (1) and (2) could be paraphrased this way:

(1r) Everyone who works hard will succeed

(2r) No one who goofs off will succeed.

This also seems to get things right for these cases, but Higginbotham argues that the analysis does not generalize. The clearest counterexamples are with the quantifiers 'most' and 'few', though there are also problem cases with all, some and none. Compare

(3) Most (of these) students will get A's if they work hard.

(3r) Most (of these) students who work hard will get A's.

(3r) will be true if a majority of the hard-working students will get A's, while (3) will be true if a majority of ***all*** the students meet the condition that they will get an A if they work hard. Neither statement entails the other. It might be that only a few students will work hard, and most of them will get A's, but that the majority of students would get B or lower, however hard they worked. On the other hand, it could be that hard work would be sufficient for an A for most of the students, but unfortunately, only a minority of those who will actually work hard are among those capable of achieving an A.

Higginbotham argues that even with universal quantifiers, there is a difference between the quantifier restrictor analysis and the quantified conditional analysis, though in this case the latter entails the former. He uses the following minimal pair to illustrate the difference:

(4) Every professor will retire early if offered a generous pension.

(4r) Every professor offered a generous pension will retire early.

"There might be," Higginnbotham says, "many professors (but even one will do) who we can be sure will not retire early, quite independently of any pension they may be offered." In this case, (4) will be false, while (4r) will

be true, if all the professors who are *actually* offered a generous pension choose to retire early.

We get the same contrast in the negative universal case, again using Higginbotham's examples:

(5) No professor will retire early if not offered a generous pension.

(5r) No professor not offered a generous pension will retire early.

"If Professor X is going to retire early, period, then he is a counterexample to [(5)], but if he is amongst those offered a generous pension, then he is no counterexample to [(5r)]."

As Higginbotham notes, to get the right result for (5), by straightforward application of compositional rules, we need to assume the principle of conditional excluded middle (CEM), which is validated by the semantics sketched above.[26] Assuming that "no x" means "for all x, not", (5) has the structure: for all x, it is not the case that [if x is not offered a generous pension, x will retire early]. CEM gets us from this to: for all x, [if x is not offered a generous pension, then x will not retire early], which seems to be the right result.

Higginbotham's discussion ignores the distinctive features of *indicative* conditionals: both the contextual constraint on the selection function used to interpret conditionals, and the prospective character of the interpretation. As he notes, his treatment of the examples "very much depends on the sensitivity of the conditional to counterfactual situations," but one cannot assume, in general, that the correctness or adequacy of an indicative conditional claim (if Professor X is offered a generous pension, he will retire early") will correspond to the truth or acceptability of the corresponding counterfactual (if Professor X, who will not in fact be offered a generous pension, *were* offered a generous pension, he would retire early.) Suppose we know that our administration will offer pensions only to those they predict will accept, and while we have no idea who will be offered a pension in exchange for early retirement, or who is disposed to accept the offer, we believe the crafty administrators are very good at such predictions. Then we might believe, of each professor, that he or she will retire early if offered a pension, while not being prepared to accept that every professor *would* accept the deal, if offered. But this hypothesis is contrived. Higginbotham's examples still work if we assume, as is plausible for this example in the normal case, that there is no epistemic or causal correlation between those who will be offered pensions in exchange for early retirement and those who are disposed to accept them. For some other examples, however, including one discussed by Higginbotham, the distinctive features of indicative conditionals matter, and ignoring them raises new problems.

Suppose (to use another of Higginbotham's examples) I have been told by someone I trust, and who shares my standards for boringness, that

(6*) every book on that shelf with a red cover is boring.

At least if we assume that we don't know anything in advance about what books are on the shelf, this seems to be enough to justify my statement

(6) every book on the shelf is boring if it has a red cover.

But, Higginbotham observes,

we know in advance that giving a book a red cover does not alter its contents, so does not affect whether it is boring. Let *b* be a book on that shelf with a blue cover. In the closest possible world, whatever it is, in which *b* has a red cover, it is boring or not, just as it is boring or not as things are. But then it seems that [(6)] should be false if there are non-boring books on the shelf whose covers are not red, although they might have been!

This is, as Higginbotham observes, the wrong result. I don't understand his response to the problem, but it seems to be an ad hoc fix. Instead, I think we should recognize what we have independent reason to recognize: that the contextual and epistemic relations that are relevant to the interpretation of indicative conditionals can come apart from the causal relations that are relevant to counterfactual conditionals.

Suppose I learn that one of the books on the shelf is *Infinite Jest* by David Foster Wallace. I know nothing about the book, let us suppose. I plan to check the cover before deciding whether to read it, since I believe (based on the information from my trustworthy informant) that if it has a red cover, it will be a boring read. The same could be said for any of the books on the shelf, assuming that I don't have an independent opinion about any of them. On the other hand, suppose I do know about *Infinite Jest* – I've read it, and it is definitely not boring. I conclude (again, based on the information from my informant) that it must not have a red cover. I don't, however, have any independent reason to have an opinion about the color of this book's cover, so if I were to learn that it has a red cover, I would conclude that the informant was mistaken. In this case (after I have learned that *Infinite Jest* is on the shelf, but before learning what color its cover is), I still believe that every book on the shelf with a red cover is boring, but I am not longer prepared to say, of every book on the shelf, that it is boring if it has a red cover.

For simple indicative conditionals, on our account, it is a constraint that the antecedent must be compatible with the CG-context. I can say "if ϕ" (indicative) only in a context that allows that it might be that ϕ. When the constraint is not met by what the addressee takes the common ground to be, he must accommodate.[27] The same constraint should be met with quantified conditionals: if I say that all or most or some of the Fs are H if they are G, then it should be true of each of the F's that it might be G. But it is not always straightforward how to apply this constraint, since even when it seems, in a sense, true of each of the Fs that it might be G, there may be ways of describing one of the Fs such that it is *not* compatible with the common ground that an F fitting that description be a G. Suppose we know that not all of the books on the shelf have red covers. It is common ground that there are some blue ones. It is compatible with this assumption that each of the books on the shelf (the first one, the second one, etc.) might have a red cover. But it is of course not true that the first ***blue*** one might have a red cover. It may be unclear, in some cases, whether it is right to say that each of the Fs might be a G, but these will be just the cases where it is unclear whether the quantified conditional is acceptable. Higginbotham mentions this example,

??(7) Every coin is silver if it is in my pocket

which is distinctly odd, in contrast with

(7*) Every coin in my pocket is silver

which is fine. It would be equally odd to say that every coin is such that it might be in my pocket.

My original discussion of the account of indicative conditionals that I am defending here began with what I called the direct argument:

Either the butler or the gardener did it, so if the butler didn't, the gardener did.

The puzzle was that if this argument, which seems compelling, is *semantically valid*, then the material conditional analysis of the indicative conditional must be right. But it seems that we have other reasons to reject that analysis.[28] The solution to the puzzle was to use the dynamic interaction of the semantic analysis with the pragmatic constraint to explain how the argument could be a compelling inference, even though not semantically valid. The acceptance of the premise changed the context in a way that ensured that the proposition expressed by the conclusion, relative to that revised context, would be true. So the semantic/pragmatic analysis accepts, and explains, the data that motivate the material conditional analysis without being saddled with the problems faced by that analysis. The possible-worlds semantics for the conditional is a part of the theory that accommodates and explains these data, but the contextual constraint, distinctive to the indicative, is also an essential part of the explanation. Barbara Abbott, in her Gricean defense of the material conditional analysis, presents counterexamples to the possible-worlds semantic analysis, but she ignores the pragmatic side of the analysis, which shows why the examples she gives are not counterexamples. Nevertheless, her main example is interesting, involving quantifiers, so let's look at it, and see how our semantic/pragmatic account explains her, and our, judgments about it.

Here is Abbott's Snodgrass example:

We have received a number of letters about the water shortage. Almost all of them were 5 pages or less, and all of those received an answer. One letter (from Byram Snodgrass) was 5 pages plus a few words, and the last letter was 8 pages. We did not reply to the last two letters. The 8-page one was just too long to consider, and Byram Snodgrass is a crank who has been writing incoherent letters to us about everything under the sun ever since we took on the post of Water Commissioner. We never answer his

letters.

Byram called our office to find out whether his letter had been sent a reply. Based on the truth in (12),

(12) Every letter no longer than 5 pages was answered.

we said (13):

(13) If your letter was no longer than 5 pages, it was answered.

Our reply was truthful.

There is a sharp contrast between the true indicative conditional in (13) and the corresponding subjunctive conditional in (14), which is not true:

(14) If your letter had been no longer than 5 pages, it would have been answered.

As noted, we never answer letters from Byram Snodgrass.

All of this seems right, and it is exactly as predicted by our analysis of the indicative conditional. The reply to Byram's inquiry was *misleading*, as I am sure Abbott would agree, since it implicated that the speaker was not in a position to give a complete answer to his question ("was my letter answered?"). But our account agrees that it is truthful, since even though both the speaker and the addressee know that the letter was longer than five pages, and so that the antecedent of the conditional is false, this is not common ground. (the indicative antecedent indicates that it is not.) On our analysis, the conditional is true in all possible worlds compatible with the prospective common ground, and we may presume that this includes the actual world, since none of the relevant presuppositions are false. The indicative conditional claim (on our modal analysis) does exactly what the assertion of the material conditional does: it excludes possible worlds in which the antecedent is true and the consequent false.

The Water Commissioner's office might have replied to Snodgrass's inquiry (equally truthfully, and equally misleadingly) with a universal generalization:

(15) Every letter was answered if it was no more than five pages long.

Both our analysis and the material conditional analysis predict that this is true. Or, the office might have replied (again, misleadingly but truthfully) with a negative universal generalization:

(16) No letter was answered if it was more than five pages long.

In this case, the reply will inform Byram (assuming he remembers how long his letter was) that his letter was not answered, but will be misleading because it implicates that the answer explains why it was not answered, and this is false. But in this case, the material conditional analysis gets the truth conditions dramatically wrong, since the negative universal material conditional falsely implies that every letter was longer than five pages, and that none of them was answered. We can save the material conditional analysis, in the negative universal case, only by giving up compositionality. But our truth-conditional version of the conditional assertion account gives a compositional analysis that I think gets the facts right.

References

1- Abbott, B. 2008. “Presuppositions and common ground,” *Linguistics and Philosophy* 31: 523-538.
2- Abbott, B. 2010. “Conditionals in English and first order predicate logic,” in D. Shu & K. Turner, eds., *Contrasting meaning in languages of the east and west.* Oxford: Peter Lang, 579-606.
3- Donnellan, K. 1966. “Reference and definite descriptions,” *Philosophical Review* 75: 281-304.
4- Edgington, D. 1986. “Do conditionals have truth-conditions?” *Critica* 18: 3-30.
5- von Fintel, K. and A. Gillies. 2011. “’Might’ made right,” in A. Egan and B. Weatherson, eds., *Epistemic modality*. Oxford: Oxford University Press, 108-130.
6- van Fraassen, B. 1966. “Singular terms, truth value gaps, and free logic. *Journal of Philosophy* 63: 481-95.
7- Higginbotham, J. 2003. “Conditionals and compositionality,” in J. Hawthorne and D. Zimmerman, *Philosophical Perspectives*, 17, *Language and Philosophical Linguisitcs*. Oxford: Blackwell, 181-194.
8- Lewis, D. 1973. *Counterfactuals*. Cambridge, MA: Harvrd University Press.
9- Lewis, D. 1974. “Semantic analysis for dyadic deontic logic,” in S. Stenlund, *Logical theory and semantic analysis Essays dedicated to Stig Kanger on his fiftieth birthday*. Dordrecht: Rediel. (Reprinted in Lewis, 2000, 5-19.
10- Lewis, D. 1975. “A problem about permission,” in E. Saarinen et al, eds. *Essays in honour of Jaakko Hintikka*, Reidel, Dordrecht, 163-75. (Reprinted in Lewis, 2000, pp. 20-33. Page references to the reprinted version.)
11- Lewis, D. 2000. *Papers in Ethics and Social Philosophy*. Cambridge: Cambridge University Press.
12- Ninan, Dilip. 2005. “Two puzzles about deontic necessity.” *New work on modality, MIT working papers in linguistics* 52. ed. by J. Gajewski, V. Hacquard, B. Nickel and S. Yalcin. Cambridge, Mass.: MIT Press.
13- Stalnaker, R. 1968. “A theory of conditionals,” in N Rescher, ed., *Studies in Logical Theory*, Oxford, Blackwell, 98-112.
14- Stalnaker, R. 1975. “Indicative conditionals,” *Philosophia* 5: ??. (Reprinted in Stalnaker, 1999, 63-77.)
15- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: The MIT Press.
16- Stalnaker, R. 2008. “A response to Abbott on presupposition and common ground” *Linguistics and Philosophy*, 31: 539-44.
17- Stalnaker, R. 2011b. “Conditional propositions and conditional assertions,” in A. Egan and B. Weatherson, eds., *Epistemic modality*. Oxford: Oxford University Press, 227-248.
18- Yablo, S. 2011. “A problem about permission and possibility,” in A. Egan and B. Weatherson, eds., *Epistemic modality*. Oxford: Oxford University Press, 270-294.

19- Yalcin, S. 2007. “Epistemic modals,” *Mind*, 116: 983-1026.
20- Yalcin, S. 2011. “Nonfactualism about epistemic modality,” in A. Egan and B. Weatherson, eds., *Epistemic modality*. Oxford: Oxford University Press, 295-332.
21- Yalcin, S. and J. Knobe. 2010. “Fat Tony might be dead. An experimental note on epistemic modals.” (http://www.semanticsarchive.net/Archive/TdiZjA3N/)

Notes

1 I am grateful to Seth Yalcin for extensive comments on an earlier draft of this paper, and email correspondence that helped me to better understand both his view and my own.

2 Yablo 2011 also consider Lewis's permission game and the comparison of deontic and epistemic modals. I am indebted to his insightful discussion.

3 I will sketch just enough of Lewis's game to bring out the lessons I want to draw from it. See Lewis, 1975 for the details.

4 Lewis does not include any formal account of assertion in his game, but one could easily add a notion of common ground, represented by a set of world-time pairs, and the standard assertion rule.

5 Lewis has a detailed discussion of the parallels between his semantics for counterfactuals and a semantics for conditional obligation in Lewis, 1973, section 5.1, and in a formal paper on dyadic deontic logic (Lewis, 1974b), but he does not mention, in the permission paper, the obvious connection between the structures proposed in those theories and the problem about permission. This is puzzling, but it may be because he was thinking of the command/permission game as a game where the spheres of permissibility do not represent some independently given value structure, but arise wholly from the speech acts of the master. It is specified, in defining the master/slave game, that in any model, before the master speaks, all is permitted. But even on this understanding of the project, the problem about permission shows that some further general conventions are needed to define a rule for permission, and it seems reasonable to think that the structures of the semantics for dyadic deontic logic would be relevant to stating such conventions.

One might hope that the structure that determines the permission rule would be the same as the structure that explains contrary-to-duty conditional requirements, one of the problems dyadic deontic logic is designed to address. But there will be some differences. Suppose the slave is commanded to carry rocks on Friday, and also (this is the contrary-to-duty command) to do penance on Sunday if he shirks his duty by not carrying rocks on Friday. That is the prior situation, but then the master gives the slave permission to take Friday off. It shouldn't follow that he still has to do penance on Sunday.

6 I have to acknowledge that my Cratchit's personality is quite different from the original, and also that he is here asking for trouble.

7 Lewis, 1975, 21.

8 von Fintel and Gillies, 2011.

9 Yalcin, 2007.

10 Wittgenstein, in his discussion of Moore's paradox in the *Philosophical Investigations* mentions an example with an epistemic modal ("Es dürft regnen, aber es regnet nicht"), but he does not note the contrast, in the case of supposition, with the version of the paradox that is explicitly about knowledge or belief. (L. Wittgenstein, 1953,192)

11 There is some discussion of reminders in Abbott, 2008, and in my response to this, Stalnaker, 2008.

12See for example Roberts, et al. 2009, and Yalcin 2011.

13 Yalcin's proposal, in Yalcin 2007, was that in the scope of propositional attitude ascriptions to a subject J, the information-state coordinate of the index shifts to the set of possible worlds that is compatible with J's attitude. So if Alice says, "Jones believes that it might rain", the statement will be true if and only if it is compatible with Jones's beliefs that's it is raining. One problem with this hypothesis is that, because of the context shift, the proposition about what Jones believes will be different from what Alice would say, or propose, if she said (or denied) that it might rain. Suppose Alice said, "Jones believes that it might rain, but I disagree – I think we can count on fine weather all day." To account for the disagreement, I think we should interpret the 'might' in the belief attribution in terms of the same contextual parameter, or an extension of it. Whether this is right or not, it is clear (as Yalcin himself brings out in later work) that his proposal does not generalize to the case of knowledge. Suppose Jones mistakenly takes himself to know that it won't rain, while in fact it will rain. Then it is clearly false that he knows that it might rain, but it is compatible with his knowledge that it will rain. I will leave questions about epistemic modals in the context of attitude ascriptions for another time, but my tentative conjecture is that they

should be interpreted relative to an extension of the information-state parameter determined by the basic context to possible worlds outside the context set (those compatible with the subject's attitudes), and not to an index shift.

14 This classic example comes from Donnellan 1966.

15 There has been much discussion in the literature on epistemic modals about so-called eavesdropper cases: situations where speakers in one context comment on epistemic modal claims made in a different context, and intuitions about these cases are mixed or uncertain. (See Yalcin and Knobe 2010 for experimental evidence about this.) Eavesdropper cases are one example of the problem of assessing the truth of an epistemic modal claim relative to possible worlds outside of the context in which that statement is made.

16 cf. Dilip Ninan, 2005, where it is observed that one cannot say things like "He must go to confession, but he won't."

17 Stalnaker, 1975.

18 Stalnaker, 1984, 111.

19 See Stalnaker 1968 and Stalnaker 1984, chapter 6.

20 Since by definition, $f(w,\phi) \in \phi$, this constraint presupposes that $\phi \cap C$ is nonempty. It is required that the antecedent of an indicative conditional be compatible with the prior context.

21 The scope of the supposition may, in some cases, continue beyond a conditional sentence.

22 This was suggested in Stalnaker, 2011, but I did not discuss there the prospective character of conditional assertions.

23 The supervaluation account of truth-value gaps was first proposed in van Fraassen 1966, and has been widely applied since.

24 Lewis, 1976, 305.

25 Higginbotham 2003.

26 CEM is the thesis that the following schema is valid: $(\phi > \psi) \vee (\phi > \sim\psi)$

27 "Presupposition accommodation" usually refers to a situation where a proposition is *added* to the common ground when appropriate speech requires it. But sometimes appropriate speech requires that a proposition *not* be implied by the common ground, and in this case accommodation requires changing the common ground so that the proposition is no longer presupposed.

28 Dorothy Edgington, in her defense of the thesis that indicative conditionals do not express propositions, also appealed to the fact that acceptance of the material conditional seems to be sufficient reason to accept the indicative conditional, at least in cases where it is an open question whether the antecedent of the conditional is true. This implies, she argued, that the indicative conditional does not express a proposition stronger than the material conditional. This, together with arguments against the material conditonal analysis, sufficed for the conclusion that indicative conditionals do not express propositions at all.