Discourse and Argument

*Draft: Please do not cite or circulate without permission.

Una Stojnić
Columbia University

April 4, 2018

1 Introduction

A simple account teaches us that argument validity is a matter of truth-preservation: roughly, an argument is valid just in case the truth of the premises guarantees the truth of the conclusion. Making this slightly more precise, we can say that a set of sentences Γ of a language $\mathcal{L}$ entails a sentence s of $\mathcal{L}$ relative to a model $\mathcal{M}$ just in case whenever all elements of Γ evaluate to truth (in $\mathcal{L}$), so does s. But this simple account becomes instantly more complex once we are dealing with inference patterns in natural languages. One immediate problem is ambiguity. For instance, consider the following:

(1) Ambiguity.
   a. If John is at the bank, he will cash a check.
   b. John is at the bank.
   c. He will cash a check.

Suppose that 'the bank' in the big premise refers to a financial institution and in the small one to the bank of a local river. In that case, the argument is clearly not truth-preserving, even though, going by the surface form, it seems to have the form of *modus ponens*. Of course, we wouldn't classify (1) as a counterexample to *modus ponens*. The argument equivocates on the ambiguity of the word 'bank', a fact that is captured by tracing the ambiguity to the difference in logical form. In particular, the lexical item corresponding to the first occurrence of 'bank' is different from the one corresponding to the second. So, an argument like (1) can be associated with distinct representations; whether an occurrence of it instantiates a particular inference pattern depends on disambiguation.

While one ought to keep track of ambiguities, this is not particularly problematic for the prospects of developing a systematic account of validity. One just ought to appreciate that the

underlying form of an argument cannot be read off of the surface form.

More complex issues arise from considerations of context-sensitivity. The fact that natural languages are riddled with context-sensitivity automatically complicates how we understand validity. For starters, consider the following inference:

(2) He is happy. Therefore, he is happy.

Is inference in (2) valid? What if we allow that the two occurrences 'he' to be resolved differently—if for instance, I use the first 'he' pointing at John, while the second one pointing at Tim? Clearly, in this case, the inference wouldn't be valid: the premise can be true while the conclusion false. But does that mean that we have a failure of the inference pattern $p \models p$? Of course, this is a wrong conclusion to draw from (2)—what's underscoring the failure is the fact that the meaning of the first occurrence of the sentence 'he is happy' in (2) is different from that of the second occurrence of that sentence, as a result of a contextual effect—a pointing gesture in this case. In a sense, it is not p that we are inferring from p, at all; it is some rather different unrelated conclusion.

A natural and common reaction to this is to say that sentences like (2) are sensitive to context, and that consequently we need to take that context-sensitivity into account in our definition of truth, entailment and validity. In particular, we need to first factor in the fact that sentences are not true or false *simpliciter*, but only relative to a context. 'He is happy' is true only relative to some context that fixes the meaning of 'he'. So, we can say that the sentence 'He is happy' is true, relative to a model $\mathcal{M}$ and a context c, just in case the object picked out by 'he' in c and $\mathcal{M}$ is happy. Since truth is now relativized to contexts, we have to correspondingly modify our notion of validity and entailment to capture this context-sensitivity. A natural and common way to do this is to hold that a set of premises entails a conclusion, just in case, for any model $\mathcal{M}$ and a context c if the premises are true in $\mathcal{M}$ and c, then so is the conclusion. This move rules out (2) as a counterexample to $p \models p$ inference pattern, as relative to the same context, 'he' has to be interpreted the same way.[1]

This idea can be seen as a ban on illicit or implicit shifts or changes in context that can affect interpretation; just as ambiguity can lead to equivocation, so can a contextual-effect of the kind

[1] One might complain that, in the case of (2), it is not the context that changes—rather, it is relative to one and the same context, but two different *demonstrations* within that context, that the two occurrences of the demonstrative select different individuals. Consequently, the two occurrences of the demonstrative are differently represented in the logical form, so that this is not an instance of $p \models p$. This view can be found in Kaplan (1989b,a). However, Kaplan himself saw this as a challenge for characterizing validity—if every two occurrences of a demonstrative come with a different directing intention, or demonstration, and are thus differently represented, we get no instances of $p \models p$ relative to *any* context. Putting this aside, note that we can make the same point with pure indexicals, which by definition cannot have more than a single value in a context. So, for instance, the validity of 'I am happy. Therefore, I am happy' shouldn't be threatened by the fact that I could have been the one who utters the first sentence, and you, the second. For this reason, Kaplan maintains, we need to characterize entailment and validity as relations between sentences-in-context, i.e. abstract pairings of contexts and sentences, rather than utterances, as real-world events. But, we will see that this move is insufficient to adequately capture validity for natural language inference patterns.

we've seen in (2).[2] The logician's ban on shifting contexts is supposed to come as a remedy to the problems raised by pernicious context-sensitivity, and rule out examples like (2) as counterexamples. The thought is simple: when one is giving an argument towards a particular conclusion, one should be committed to a stable context, one that guarantees a non-equivocal interpretation to the context-dependent expressions.[3]

2 Complications in a Stable Context

However, there are problems even if we assume a ban of implicit shifts in context. Let us take an example that dramatizes the problems. Famously, McGee (1985) argued that *modus ponens* is invalid for English indicative conditional. The counterexample goes as follows:

(3) Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third. Those apprised of the poll results believed (3-a) and (3-b), with good reason. Yet they did not have reason to believe (3-c).

a. If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.
b. A Republican will win the election.
c. If it's not Reagan who wins, it will be Anderson.

But, of course, (3-c) follows from (3-a) and (3-b) by *modus ponens*. What gives? Insofar as we would like to reject (3-c), if *modus ponens* is valid, then either (3-a) or (3-b) has to be false. It seems that denying (3-b) is a non-starter. That leaves us with (3-a). But (3-a) is intuitively true.[4]

Since we are dealing with a *modus ponens* inference, we need to say something about the truth-conditions for an indicative conditional in a context. Taking for a moment the standard account

[2] One might want to reserve the term 'equivocation' for genuine instances of ambiguity, where an argument is trading on a non-uniform disambiguation of a term, as in (1). In the case of (2), the argument doesn't involve two different *meanings* for the pronoun 'he', merely two different semantic contents for a single lexical entry. Nothing hinges on the choice of terminology. Clearly, one ought to control for both types of effects—disambiguation and context-sensitivity resolution—in evaluating an argument.

[3] This sentiment is echoed more or less explicitly throughout the literature. See in particular Kaplan (1989b,a); von Fintel (2001); McGee (1985), *inter alia*.

[4] To take a similar example, consider the following:

(i) Gentle Murderer:
a. If John will kill Bill, he ought to kill him gently.
b. John will kill bill.
c. John ought to kill Bill gently.

Something has gone wrong again. While we might be persuaded by (i-a) on the grounds that a lesser evil is better then a greater one, (i-c) certainly sounds false. And while holding the context fixed can guarantee that the premise is true when the conclusion is, this doesn't explain why (i) feels like a counterexample.

of conditionals, let us suppose that a conditional $if(p)(q)$ is true at a world w, in the context c just in case all the closest-to-w p-worlds are q-worlds (where closeness is determined by context).[5] Now, suppose that (3-c) is false in w in c. Then, it's not the case that all of the closest worlds from w in which Reagan loses, are such that Anderson wins. But is then (3-a) true at w in c? Well, supposing that (3-b) is true at w in c, then w is a world in which a Republican wins. Assuming that the w is among the worlds closest to itself, since (3-c) is false, not all the closest worlds in which a Republican wins are such that in it, it is true that if it's not Reagan, it's Anderson; w is a counterexample. So, if both (3-a) and (3-b) hold in a context, (3-c) has to hold in that context as well: Stalnaker/Lewis account validates *modus ponens*. But, McGee (1985) argues, this is a problematic result as we have good reasons to believe (3-a) and not (3-c); (3-c) seems false while (3-a) sounds true.[6] So, if relative to a fixed context (3) is valid, we still need a further explanation of why we feel that something has gone wrong.

To make things worse for *modus ponens* inference, English indicative conditional seems to obey the so-called import-export (IE) pattern, the argument illustrated in (4):

(4) There was a murder at a mansion. One of the staff members is guilty.

a. If the butler is innocent, and the gardener is innocent, the cook is guilty.

b. If the butler is innocent, then if the gardener is innocent, the cook is guilty.

(4-a) seems to entail (4-b). Indeed the two conditionals sound equivalent: one is true just in case the other is. So, *prima facie* (4) seems to conform to the so called import-export inference pattern: $p \wedge q \rightarrow r \dashv\models p \rightarrow (q \rightarrow r)$. The problem, of course, is that as a well known result by Gibbard (1981) establishes, any conditional that supports both IE and MP and is at least as strong as material implication is equivalent to material implication. Since material implication is a non-starter for the meaning of an English indicative,[7] at least one of the patterns has to go. Problematically, the standard semantics validates *modus ponens* and gives up IE. But this seems as exactly the opposite from what we observe: while we appear to have counterexamples to MP, IE is flawless.

[5] More formally:

(i) $[\![if(p)(q)]\!]^{\langle w,c\rangle} = 1$ iff $\forall w' \in Closest(w,p) : w' \in [\![q]\!]^{w,c}$, where $Closest(w,p) = \{w' | wRw'$ & $w' \in [\![p]\!]^{w,c}$ & for all $w'' s.t. wRw''$ & $w'' \in [\![p]\!]^{w,c}, w' \leq_w w''\}$

For simplicity we drop Uniqueness Assumption of Stalnaker's framework, and the Limit Assumption of Lewis' framework; this is of no significance for our purposes. Both Lewis' and Stalnaker's analysis validate *modus ponens* (assuming at least weak centering). More generally, the specifics of the account of conditional do not really matter for the counterexample so long as the account maintains that the domain of worlds over which the conditional is ranging is held constant across the argument in (3).

[6] Indeed, McGee (1985) takes (3) to be a counterexample to Stalnaker's theory of conditionals, precisely because it validates *modus ponens*, while (3) shows that *modus ponens* is invalid.

[7] Despite some venerable attempts at defending the material conditional analysis (Grice, 1989), the analysis seems hopeless (Lewis, 1975).

3 Uniformity

So given the standard account of conditionals, even when subscribing to the ban on context-shifting, we get puzzling results, seemingly precisely the opposite from intuitively correct ones. But, this result was underwritten by the type of context-sensitivity and the relevant contextual parameters the standard account employs. However, there are good reasons to think of the relevant context-sensitivity involved in the interpretation of conditional expressions—and consequently the way of controlling for a contextually uniform interpretation of these expressions—in a way that's different from what the standard account suggests.

The basic idea is simple: conditionals, and modal expressions more generally, exhibit a similar type of context-sensitivity as pronouns do (Stone, 1997, 1999; Bittner, 2001, 2014, 2011; Brasoveanu, 2010; Stojnić, 2016c,a). They require some salient possibility from the context that restricts their domain of quantification. So, just as a uniform interpretation of pronouns requires that they be interpreted relative to the same antecedent (provided from the non-linguistic context, or introduced by the prior discourse), similarly, a uniform interpretation of conditionals (and other modal expressions) requires a shared antecedent. Just as a failure to track how a pronoun is resolved can result in equivocation, so too can a failure to track how modal and conditional expressions are resolved.

That modal expressions exhibit anaphoric behavior similar to that of pronouns has been argued for in the literature (Stone, 1997, 1999; Bittner, 2001, 2014, 2011; Brasoveanu, 2010; Stojnić, 2016c,a). If we understand modal expressions as quantifiers over possibilities, where the domain of quantification is restricted by a contextually provided possibility (a restrictor), they can be restricted by a prominent possibility either from a non-linguistic context, or from a prior discourse:

(5) (Looking at high-end stereo equipment): My neighbors would kill me. (Stone, 1997)

(6) A wolf might walk in. It would eat you. (Roberts, 1989)

The same goes for conditional: they, too, are interpreted against some prominent body of information—they don't just range over all (or even just all the closest) possible worlds. And here, too, the relevant possibility can be provided by the non-linguistic context, or the prior discourse. In the same situation in which (5) was uttered, one can utter (7):

(7) If I turned the volume up, my neighbors would kill me.

And similarly the relevant restricting possibility can be given by a prior discourse. Consider the following example:[8]

[8] A similar example from Stojnić (2016c) illustrates the same point:

(i) If a wolf walks in, we will use the tranquilizer gun. If we manage to shoot it, we will be safe.

(8) (The antidote is deadly, except when administered after a snake bite, in which case it is life saving. Looking at John in front of a snake tank, I say:)
If a snake escapes, it will bite John. But, if he takes the antidote, he will live.

This in turn gives us another way of thinking of the requirement for a uniform interpretation of these expressions: just as two (or more) pronouns need to have the same antecedent (either from the prior discourse or from the non-linguistic context) in order to be interpreted uniformly, so too, in the case of modals and conditionals, they have to be interpreted relative to the same restrictor possibility.

To make this concrete, following Stojnić (2016c,a), we can represent the truth-conditional meaning for a conditional that captures the anaphoric dependency on contextually specified restriction as follows:[9]

Definition 3.1.

$Cond(p,\ q,\ r) := \{w \mid \forall w' : wRw'$, if $w' \in p$ & $w' \in q$, then $w' \in r\}$

Where p, is the anaphorically retrieved restrictor—the proposition restricting the interpretation of the conditional, q corresponds to the proposition expressed by the antecedent, and r to the one expressed by the consequent, an utterance of a conditional expresses truth-conditions corresponding to a set of worlds such that for each w in the set, all the worlds R-accessible from w, that are p and q worlds, are r worlds as well. So, an utterance of a conditional is true in w if and only if all the p and q worlds in the domain of quantification are r worlds as well. Crucially, the conditional is always evaluated against some prominent body of information, that need not correspond to the unrestricted set of epistemically live worlds discourse initially.

Similarly, we interpret modal expressions as restricted in a context by a prominent proposition. So for instance, where '$M(p, q)$' stands for the truth-condition expressed by an utterance of '*might* ϕ', q is the proposition expressed by the prejacent ϕ of the utterance of '*might* ϕ', p the proposition corresponding to an anaphorically retrieved restrictor, and 'R' denotes a contextually provided accessibility relation, that determines a set of epistemically accessible worlds from a given world w, we can define the truth-condition expressed by an utterance of '*might* ϕ', '$M(p, q)$', as follows:

Definition 3.2.

$M(p,\ q) = \{w \mid \exists w' : wRw'$ & $w' \in p$ & $w' \in q\}$

3.2 says that an utterance of '*might* ϕ' is true at w relative to a context c if and only if there are worlds R-accessible from w, in which both the anaphorically retrieved restrictor, p, and the prejacent proposition, q, are true. (*Mutatis mutandis*, we can define the truth-conditions for other

[9] For the moment, we shall ignore the ordering source—the contextually supplied parameter that ranks worlds relative to their proximity to the world of evaluation given some standard. We can easily factor this parameter back in.

modal expressions in the same way. So, for instance, we can treat 'must' as a universal dual of 'might'. The important aspect, for our purposes is the contextually supplied restriction on the domain of quantification.)

So, a uniform interpretation of a conditional across a discourse requires that it be restricted by the same restrictor throughout, just as a uniform interpretation of a pronoun requires that it receives the same interpretation throughout. To illustrate this idea consider the following example, from Gibbard (1981):

(9) Sly Pete and Mr. Stone are playing poker on a Mississippi riverboat. It is now up to Pete to call or fold. My henchman Zack sees Stone's hand, which is quite good, and signals its content to Pete. My henchman Jack sees both hands, and sees that Pete's hand is rather low, so that Stone's is the winning hand. At this point, the room is cleared. A few minutes later, Zack slips me a note which says (9-a), and Jack that which says (9-b):

a. If Pete called, he won.

b. If Pete called, he lost.

Notice that we can naturally understand the two conditionals as restricted by different bodies of information here (the first one referencing the information available to Zack, and the second, that available to Jack). So, the two conditionals are interpreted differently: the first expresses the content that *given Zack's evidence*, Pete won if he called, and the second that *given Jack's evidence*, Pete lost if he called.[10] Notice that, had Jack instead reported that 'If Pete called, he won', intuitively he would be saying something false: given his evidence, it is not the case that Pete won if he called.

So similarly, then, when we are interpreting (3) we need to control for the interpretation of the relevant conditionals. If an argument requires a stable context and a uniform interpretation, we need to ask what would a uniform interpretation amount to? Let us consider the example again:

(3) Opinion polls taken just before the 1980 election showed the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson,a distant third. Those apprised of the poll results believed (3-a) and (3-b), with good reason. Yet they did not have reason to believe (3-c).

a. If a Republican wins the election, then if it's not Reagan who wins it will be Anderson.

b. A Republican will win the election.

c. If it's not Reagan who wins, it will be Anderson.

[10] A similar point is argued for in Kratzer (1983); however, Kratzer takes the restriction to be restricting a covert modal operator in the logical form of a conditional. I make no commitment to covert modal operators in the LF. The difference between the present account and one that maintain that conditional antecedents obligatorily restrict (overt or covert) modals in the consequent will become clear shortly. Note that, as demonstrated by (8), even if the antecedent supplies a restriction, a further restriction from the context is needed. The same holds for (9-a) and (9-b).

Now, presumably the first conditional is interpreted relative to the body of information that comprises the information given in the scenario, describing the election situation, the polls, etc. But, it also embeds a conditional, namely, "If it's not Reagan, then it will be Anderson". As conditionals require to be restricted by some body of information, we then have to ask which restrictor restricts *this* conditional. More importantly, whichever restriction is in play, it has to be the same one that is in play in the interpretation of the conditional in the conclusion. If it weren't, we would get the same shift in meaning as we would were we to interpret (9-a) in the mouth of Jack in (9). Similarly, notice that such a shift in interpretation in (3) is analogous to the change in interpretation between we get between (10-a) and (10-c) in (10):

(10) a. If John is at home, he is reading a book.
b. John is at home.
c. (Pointing at Tim) He is reading a book.

While allowing the two occurrences of the pronoun to resolve differently clearly yields an argument which is invalid, given that the interpretation of the pronoun isn't uniform the pattern isn't really a counterexample to *modus ponens*. Indeed, one naturally concludes that in (10), precisely the kind of illicit context-shift is taking place that was banned by the ban on context-shifting.

So, given that conditionals require a contextually supplied restrictor, then just as in the case of (10), in order for (3) to exemplify an instance of *modus ponens*, the interpretation of the corresponding conditionals has to be uniform. In other words where p is some prominent body of information, the conditional 'if it's not Reagan, then it will be Anderson', has to be interpreted as 'if p and it is not Regan, it will be Anderson'. Somewhat more precisely, where the truth condition for our conditional is as in 3.1, p is some prominent contextually provided restrictor, q the proposition that Reagan didn't win, and r the one that Anderson won, the conditional in (3-c) and the one in the consequent of (3-a) have to express the same proposition, namely:

(11) $\{w \mid \forall w' : wRw', \text{ if } w' \in p \ \& \ w' \in q, \text{ then } w' \in r\}$

Of course, there is a substantive question of which body of information p is. For instance, suppose this body of information is one that includes the information that a Republican will be the winner; in that case, both the first premise, and the conclusion will be true. Or suppose it was one that is not opinionated on the winner's party. In that case, both the first premise and the conclusion would be false. In both cases we have an instance of *modus ponens*. In one, we transition from true premises to a true conclusion; in the other, one of the premises is false, and so is the conclusion. This is as it should be.

Of course, once we start interpreting the two conditionals non-uniformly, say, interpreting the one in the consequent of (3-a) as restricted by a body of information in which it's decided that

a Republican will win, while the one in (3-c) as restricted by a body of information that is unopinionated in this regard, we get a case where the premises are true, and the conclusion is false. But this is a case where there is a contextual shift between premises, one akin to the one in (10). And of course, *that* is not a threat to *modus ponens*.[11]

However, this is not all there is to be said about this case. If we understand the requirement for a uniform interpretation of context-sensitive items as a ban on context-shifting, then we seem to face a problem in (3). Suppose the context fixes a particular body of information—a particular prominent possibility—as the contextual parameter relevant for modals and conditional expressions. But then, we would expect that not only the conditional in the consequent of (3-a), and the one in the conclusion in (3-c), but also the embedding conditional in (3-a) be restricted by it. But this now raises an important problem: (3) has no natural interpretation on which the three conditionals receive a uniform interpretation! Specifically, while the embedding conditional in (3-a) naturally is interpreted against a body of information unopinionated with regards to the winner's party, the embedded conditional has no such natural interpretation: it has no reading on which it simply means that *given all available information and regardless of the winner's party affiliation* the winner will be Anderson if it isn't Reagan.[12]

An account that simply requires that the context is held constant and the interpretation uni-

[11] Similar point about the uniformity of information is made in, e.g., Gillies (2010). However, Gillies takes this kind of data to call for a revision of the classical notion of consequence, as well as a challenge to the standard truth-conditional account of meaning. Gillies offers an alternative account of the interaction between context and content, one quite different than the one I offer here. I shall return to the differences between the two accounts, and argue that the account I favor is more adequate in § 6.

[12] Of course, many accounts build this dependence into the meaning of the conditional, either by positing that a conditional antecedent restricts a (covert or overt) modal operator (see Kratzer (1983)), or that it shifts the relevant body of information against which the conditional is interpreted (see Veltman (1985); Gillies (2010, 2004); Yalcin (2007), *inter alia*). While the present account will predict such restriction in most cases, it doesn't force it. This seems as a welcome prediction, as sometimes, indeed, conditionals don't seem to restrict modal operators in their consequent:

(i) If Britney Spears drinks coke in public, then she ought to. (Zvolenszky, 2002).

Second, even if antecedents of conditionals obligatorily restricted the interpretation of modalized items in their consequents, this does not deliver a solution general enough. For instance, it doesn't predict the following (and see Stojnić (2016c) for a more extensive discussion of this issue):

(ii)

a. Suppose the winner is a Republican.
b. Then if it's not Reagan, it's Anderson.
c. The winner will be a Republican.
d. So, if it is not Reagan, it's Anderson.

Finally, even if conditionals obligatorily restricted the interpretation of the modalized expressions in the consequent, that observation could be captured in a variety of different accounts, including the present one which posits an anaphoric element in the interpretations of the conditionals. But, as I shall argue, the present account has advantages in capturing and individuating inference patterns expressed by natural language discourses.

form, doesn't explain this phenomenon: indeed, an account that simply requires that all the relevant context-sensitive expressions be resolved uniformly predicts precisely the reading that seems impossible to get. Is there a way of capturing the requirement of uniform interpretation of component context-sensitive expressions in examples like (3), while at the same time doing justice to this observation? Indeed, I will suggest, there is. What we need is an account of how the context affects the interpretation of context-sensitive items, in a way that allows us to control for the uniformity that a particular inference pattern requires, while not simply requiring that all the relevant context-sensitive components an argument contains be interpreted in a coordinated way. The insight builds on Stojnić (2016c,a), who argues that the resolution of the contextual restriction on the interpretation of modal and conditional expressions is determined by discourse internal mechanisms that organize discourses into coherent units. I sketch the account presently.

4 Discourse Structure and the Structure of Argumentation

The insight behind the account developed in Stojnić (2016c,a) builds on the observation that discourses are more than a sum total of their individual sentences taken in no particular order. An example that nicely illustrates the point is the following one from Hobbs (1979):

(12) John took the train from Paris to Istanbul. He has family there.

(13) John took the train from Paris to Istanbul. He likes spinach.

(12) isn't just a random set of sentences about John; the discourse conveys that John took a train from Paris to Istanbul *because* he has family there. Understanding that there is an explanatory connection between the two bits of discourse is part of understanding the overall contribution of (12). But the requirement for a meaningful connection carries over to the oddly sounding (13). We are left wondering how the preference for spinach explains the train trip to Istanbul. The failure to establish a connection leads to a failure to understand the discourse.

Coherence theory proposes to capture this insight by positing an implicit organization to the discourse, a network of coherence relations that signal how individual utterances are connected into a coherent whole (Hobbs, 1979, 1990; Kehler, 2002; Asher and Lascarides, 2003). So, (12) harbors a coherence relation of Explanation signaling that the second sentence explains the event described by the first one; similarly, in (13) one expects that the second sentence stands in Explanation relation to the first one. As the relation isn't established—we are wondering how does the spinach preference explain the train trip?—the discourse is infelicitous.[13]

[13]Note that, of course, with sufficient background information the discourse can be made felicitous: e.g., suppose we learned that Istanbul is famous for its fine spinach. This is just as expected on the coherence theoretic account: the point is precisely that, absent background information the relation cannot be confirmed. Note also that it's important that, out of the blue, we are trying to confirm the relation of Explanation, even though a number of

What is of concern for our purposes is the observation that coherence structure of a discourse affects the resolution of context-sensitive items, e.g., of demonstrative pronouns. This is illustrated by the following example:

(14) Phil tickled Stanley. Liz poked him. (Smyth, 1994)

In (14) the resolution of the pronoun 'him' co-varies with the choice of the organizing rhetorical relation. If the relation is that of Result, where the second sentence describes the result of the eventuality described by the first one, the pronoun is resolved to the subject of the previous clause 'him'. If it is organized by Parallel, comparing the two eventualities according to their (dis)similarities, the pronoun is resolved to the antecedent introduced by the NP in the same grammatical position, that is, Stanley.

The observation that discourse structure affects the resolution of demonstrative pronouns is well supported by empirical research (Wolf, Gibson, and Desmet, 2004; Kehler et al., 2008). Stojnić, Stone, and Lepore (2017) argue that these effects of coherence relations on pronoun resolution are conventionalized: as a matter of their linguistic contribution, coherence relations make a particular candidate referent prominent, and pronouns as a matter of their meaning select the most prominent candidate referent which fits their associated character.[14] Somewhat more precisely, Stojnić, Stone, and Lepore (2017, 2013) propose to model this effect of discourse coherence as an effect on context, conceived of as a conversational scoreboard, a la Lewis (1979), a running record of contextual parameters. The context keeps track of how parameters change with an evolving discourse. It includes, among other parameters a ranking of candidate referents according to their relative prominence—those higher in the ranking being more prominent than those lower.[15] Given this setup, coherence relations are modeled as associated with a grammatically specified update to this ranking, making certain referents prominent, demoting others. Pronouns then as a matter of their meaning select the highest ranked referent that satisfies their character: roughly, 'he' selects the

other relations would fit (12) well: e.g., it could be taken as a description of John's odd, or old-fashioned tastes, where the discourse would be organized by a resemblance relation Parallel (cf. Lepore and Stone (2015)). That a particular relation is preferred out of the blue suggests that it is primed by discourse internal cues—e.g. by the thematic structure of the verb in the first clause. (Cf. Stojnić, Stone, and Lepore (forthcoming); Lepore and Stone (2015); Kehler (2002); Asher and Lascarides (2003). For a discussion of the typology of rhetorical relations see Kehler (2002); Knott (1996); Asher and Lascarides (2003); Mann and Thompson (1988).)

[14] Specifically, Stojnić, Stone, and Lepore (2017) argue that affects of coherence relations on the resolution of pronouns cross-cut pragmatic considerations and abductive reasoning strategies that aim to establish the overall most plausible interpretation. Moreover they point out cross-linguistic variations in associations between coherence relations and prominence marking, which suggest that these effects are manifestations of learned rules of a particular language, rather than by-products of common-sense reasoning. In the interest of space, I do not rehearse these arguments here.

[15] One can think of this context as a kind of an extension of the Kaplanean notion of context, taking into account how it *changes* across a span of a discourse. Note that, officially, a Kaplanean context separates out the assignment function from the context proper, but we can think of parameters of the Kaplanean context as elements of contextual assignments. See Stojnić, Stone, and Lepore (2017).

top-ranked element in the ranking that is a male.[16]

Importantly, discourse structure doesn't merely affect the interpretation of a pronoun: it affects the interpretation of context-sensitive expressions more broadly (Stojnić, 2016b), and in particular the interpretation of modal expressions and conditionals (Stojnić, 2016c,b). As argued by Stojnić (2016c,b), just as in the case of pronouns, in the case of modals and conditionals, discourse structure dictates the prominence of possibilities that are candidate restrictors for the resolution of subsequent modal anaphora. We see this already with simple cases of modal subordination, as the one in (6), repeated here:

(6) A wolf might walk in. It would eat you. (Roberts, 1989)

The first sentence introduces a hypothetical possibility, and the second further elaborates on it. As Stojnić (2016c,a) argues, Elaboration relation makes this possibility prominent, and consequently the modal in the second sentence selects it as its restrictor. The same effect we observe in (15):

(15) If a wolf walks in, it would eat you.

The antecedent introduces a hypothetical scenario, and the consequent elaborates on it. The Elaboration relation makes the possibility elaborated on prominent, and the modal 'would' in the consequent selects it as a restrictor.

These effects of discourse structure on the interpretation of conditionals is general: the interpretation of sequences of conditionals across a discourse crucially relies on the choice of a coherence relation. For instance, compare the following:

(16) If the snake escapes, it will bite you. If you get the antidote, you will live.

(17) If the snake escapes, it will bite you. If the bear escapes, it will hunt you down.

In the first example, the second conditional further elaborates on the possibility introduced by the first, in which the snake escapes the cage, and bites the addressee. In turn, its interpretation is restricted by this possibility: in case the snake escapes and bites you, given that you get an antidote, you will live. On the other hand, in (17), the two conditionals are organized around Contrast relation, comparing the possibility in which a snake escapes with one in which the bear escapes. In this case we have a different interpretation. Contrast requires that the two bits of discourse provide contrasting information about some body of information—e.g., about the body of information comprising what's taken for granted for the purposes of the conversation—and in turn this body of information is made prominent and is what restricts the interpretation of the second

[16]For a formal implementation of this idea see Stojnić, Stone, and Lepore (2017). It is worth noting that the account allows us to capture not just referential, but also bound and discourse bound (e-type) readings of pronouns, while assigning each demonstrative pronoun a uniform, unambiguous linguistic meaning and representation.

conditional: so, we get an interpretation that given what we take for granted, if the bear escapes, it will hunt the addressee down.[17]

As in the case of pronouns, Stojnić (2016c,a) argues that these effects are conventionally encoded as a linguistic contribution of coherence relations. To illustrate, consider the following:

(18) If it is not raining and it might be raining, then I am uninformed about the current weather. (Yalcin, 2007)

As Yalcin notes, strangely, (18) seems incoherent. This is strange on the standard accounts of the meaning of 'might', as 'it might be raining' should just carry the information that the contextually prominent body of information i is compatible with raining. Of course, for many values of i, then, (18) should have a perfectly consistent, even true reading. Yet, it seems it does not. Stojnić (2016a) argues that this is due to a conventional effect of discourse coherence on the interpretation of the modal 'might'. The first conjunct of the conjunction in the antecedent of (18) introduces a hypothetical scenario in which it is not raining. The second conjunct elaborates on this scenario, and much as in (6) or (15), Elaboration makes this scenario prominent. As a result, the modal in the second conjunct is restricted by it. This results in an inconsistent interpretation: *it is not raining and it might be given that it is not.* The important point to note is that, were the effect of Elaboration a mere defeasible byproduct of pragmatic reasoning, one would expect that (18) should have a perfectly consistent interpretation: indeed, such interpretation should not only be available but preferred, as plausibility, charity, and relevance considerations would all favor it. That (18) has no such interpretation suggest that the effect of Elaboration is a matter of a linguistic rule.[18]

Extending the picture of context already sketched, Stojnić (2016c,a) proposes to model these effects on the interpretation of modal and conditional anaphora in the similar fashion as in the case of pronouns. Just as the context, understood as a conversational scoreboard, included a ranking of relative prominence of candidate referents for the resolution of pronouns, so now we add to it a ranking of candidate possibilities that can serve as restrictors for the domain of quantification of modal and conditional expressions. As in the case of pronouns, the ranking changes with the evolving discourse, and is affected by linguistic mechanisms that introduce and promote possibilities—antecedents of conditionals, modal expressions, coherence relations.

With the proposal just sketched in mind, we can now expand on our prior explanation of the example in (3). (3-a)–(3-c), let us suppose, is uttered against some prominent body of epistemically live worlds: presumably the one which contains the information about the election provided by (3). Relative to this body of information the antecedent of the big, conditional premise, (3-a), introduces a possibility in which a Republican is the winner. The consequent of this conditional

[17] For a detailed discussion, and further examples, see Stojnić (2016a,c).

[18] See Stojnić (2016a) for a detailed discussion of this point, as well as further arguments for treating the effects of coherence as a matter of linguistic convention.

further elaborates on this scenario. The Elaboration relation between the antecedent and the conditional in the consequent makes the possibility elaborated on prominent, and this possibility in turn serves as the restrictor for the conditional in the consequent. Relative to this possibility, the one comprising epistemically live worlds in which a Republican is the winner, the antecedent of the embedded conditional introduces a further new possibility, the one in which the winner is not Reagan. The consequent of the embedded conditional further elaborates on this possibility, and Elaboration relation makes this possibility—the one comprising epistemically live worlds in which a Republican, but not Reagan wins—prominent. Finally, the consequent of the embedded conditional provides further elaboration about this possibility: in it the winner is Anderson. Given our truth condition in 3.1, (3-a) is true, then, just in case, all the epistemically live worlds in which a Republican other than Reagan wins are the ones in which the winner is Anderson. The second premise further elaborates on the information provided by the first one—on the body of live possibilities in which the conditional proposition expressed by the first premise holds, requiring further that throughout this possibility, a Republican is the winner. Finally, the conclusion is related to the premises by relation of Conclusion. Conclusion promotes the body of information incorporating the information expressed by the premises—in particular, this will be the body of information comprising live epistemic scenarios in which the winer is a Republican. This body of information in turn restricts the interpretation of the conditional in the conclusion: relative to it, the antecedent introduces a hypothetical scenario, one in which Reagan loses. The consequent further elaborates on this scenario, requiring that in it the winner is Anderson. The conditional in the conclusion is then true just in case within all the epistemically live worlds, in which a Republican wins, and in which Reagan loses, Anderson is the winner. Note that this is precisely the content expressed by the consequent of the conditional premise, and precisely the content that MP-inference commits us to.

Notice that the account doesn't require a single context to be held fixed in interpreting (3). On the contrary, it predicts that the context has to change as a result of linguistically induced effects of its component expressions, and its underlying structure: the antecedents of corresponding conditionals, and organizing coherence relations thread the change in prominence in a way that determines the propositional content expressed (Stojnić, 2016c). But precisely in virtue of the change in context the argument ends up expressing the underlying pattern corresponding to *modus ponens* inference: in particular, it guarantees that the conditional in the consequent of the conditional premise, and the one in the conclusion are interpreted relative to the same body of information, and hence express the same content. Notice, also, that we can explain why (3) seems troubling in the first place: considering (3-a)—(3-c) as individual sentences, rather than a single pointed piece of discourse, we then have to first determine some body of information that restricts (3-c).[19] The

[19] A similar point is made in Gillies (2004), who points out that one cannot simply evaluate (3-c) in isolation any more than one can evaluate 'He sat down' without providing a context that fixes the meaning of 'he'. However,

intuition that we should reject (3-c) is precisely underwritten by interpreting it against a body of information that is unopinionated with respect to the winner's party affiliation. However, as soon as we understand (3-c) as a conclusion drawn from the premises that precede it, we get a different interpretation of the conditional, the one that fits the *modus poenens* pattern. This is clear once we consider the impeccable (19) in contrast to odd sounding (20):

(19) A Republican will win, so if it's not Reagan, it will be Anderson.

(20) A Republican will win, so if it's not Reagan, it will be Carter.

Notice that this account also allows us to capture (4), repeated below:

(4) There was a murder at a mansion. One of the staff members is guilty.

a. If the butler is innocent, and the gardener is innocent, the cook is guilty.

b. If the butler is innocent, then if the gardener is innocent, the cook is guilty.

Since on our account conditionals are interpreted relative to some prominent body of information, we will have to determine which bodies of information are relevant for the interpretation of conditionals in (4-a) and (4-b). Plausibly, (4-a) is uttered against the body of information described in the scenario, the one in which we know that one of the staff members is the culprit. The antecedent of the conditional in (4-a) contains a conjunction. The first conjunct introduces a hypothetical scenario in which the butler is innocent, and the second one further elaborates on it, requiring that it also be a scenario in which the gardener didn't do it, either. The consequent further elaborates on this scenario, requiring that in it, the cook is the culprit. So (4-a) will be true just in case all the epistemically accessible worlds in which a stuff member is guilty, and in which the butler and the gardener are both innocent are ones in which the cook is guilty, i.e. just in case, given all we know, and given that the butler and the gardener are innocent, the cook is the culprit.

Now, take (4-b). Relative to the same information described in the scenario, the antecedent introduces a hypothetical scenario in which the butler is not the culprit. The conditional in the consequent further elaborates on this scenario, and Elaboration relation makes this scenario prominent. The conditional in the consequent, as a result, is restricted by it. Then, relative to this body of information, the antecedent of the embedded conditional introduces a hypothetical scenario, one in which the gardener is innocent. The consequent further elaborates on this scenario. The Elaboration relation makes the scenario prominent, and the consequent requires that it be such that in it, the cook is the culprit. Thus, (4-b) is true just in case all the epistemically accessible worlds in which a stuff member is guilty, and in which the butler and the gardener are both innocent are

Gillies uses this to motivate an account that doesn't ultimately treat conditionals as anaphoric expressions, nor as expressing non-trivial propositional content. I will return to the comparison between such an account and the one endorsed here in what follows.

ones in which the cook is guilty, i.e. just in case, given all we know, and given that the butler and the gardener are innocent, the cook is guilty. So we see that given our truth-conditions in 3.1 and given our account of context-sensitivity resolution, (4-a) and (4-b) express the same content. However, notice also that, due to the contextual restriction on the interpretation of the conditional in the consequent of (4-a), this embedded conditional simply expresses the same truth-condition as (4-b): given all the epistemically live words in which neither the butler nor the gardener is guilty, the cook is. Of course, if the resolution of the conditional in (4-b) weren't restricted in this way—if the embedded conditional could range over any body of information in the context, then of course, then of course exportation would not be truth-preserving.[20]

Notice that requiring the conditionals in (4) to be uniformly interpreted—requiring, that is, that (4-a), (4-b) *and* the conditional in the consequent of (4-b) all be interpreted relative to the same body of information—would allow for counterexamples to exportation. It is only if we require the specific type of contextual shift, on the present account induced by Elaboration relation, that we can guarantee that (4-a) will entail (4-b). But since on our account Elaboration is associated with a linguistically mandated effect on the interpretation of a subsequent expression, does that mean that our semantics validates exportation after all? This raises a question of how to individuate inference patterns in the presence of context-sensitivity.

5 The Structure of an Argument

If we merely required that all the relevant context-sensitive expressions receive uniform interpretation throughout an argument, then why certain interpretations aren't witnessed becomes mysterious. For instance, suppose we represent the uniform interpretation by co-indexing the relevant contextually sensitive items. Then, we can't really explain why we don't witness the pattern we see below:

(21)

a. If$_i$ a Republican wins, then if$_i$ the winner is not Reagan, it'll be Anderson.

b. A Republican will win.

[20]Suppose that the conditional in the consequent of (4-b) can just select our overall body of knowledge. Then (4-a) could be true while (4-b) false. Suppose that (4-b) is true at a world w. Then all w' accessible from w among the worlds compatible with what we know, in which the butler and the gardener are innocent, have to be such that in them the cook is guilty. But, it could still be false that all the worlds w' accessible from w among the worlds compatible with what we know, in which the butler is innocent, are worlds in which it is true that *given all we know*, if the gardener is innocent the cook is guilty. After all, we might not know in w' that the butler is innocent. This isn't surprising: we wouldn't expect exportation to hold for a strict conditional.

It's simple to see the formal point. Assume 3.1. Take a model with two worlds, w_1, w_2. Let E be some restrictor body of information, and take some p, q, r such that p, E, q, r all hold in w_1 and $\neg p, E, q, \neg r$ hold in w_2. Suppose that R is reflexive and that $w_1 R w_2$. Then, $Cond(E, p \wedge q, r)$ holds at w_1, but $Cond(E, p, Cond(E, q, r))$ doesn't. Namely, $\forall w', w_1 R w'$, if $w' \in E \wedge p \wedge q$, then $w' \in r$, but there is a w', namely w_1, such that $w_1 R w'$ and $w' \in E \wedge p$ and $w' \notin \{v | \forall v' \text{ s.t. } v R v', \text{ if } v' \in E \wedge q, \text{ then } v' \in r\}$.

c. So, if$_i$ it is not Reagan it will be Anderson.

And similarly, we cannot explain why it is that we do not see the pattern that we see below:

(22) a. If$_i$ the butler didn't do it and the gardener didn't do it, then it was the cook.
b. If$_i$ the butler didn't do it, then if$_i$ the gardener didn't do it, it was the cook.

Now, we have seen that we can tell a systematic story of how the context changes in a way that governs the interpretation of context-sensitive items which, in turn, allows us to capture the patterns we witness. The story required that we take into account how the argument is structured into a coherent piece of discourse, and not merely regard it as a set of premises and conclusions. But if the structure matters, and if it governs the resolution of context-sensitivity by manipulating context, then the question arises as to the validity of the above inference patterns. After all, our initial notion of validity required a fixed context throughout the evaluation of an argument. If we want to capture the idea that the context can shift as affected by linguistic components of an utterance, then the question becomes what is it for an argument to be valid in such a shifty environment, and how do we individuate the logical form of an argument in the presence of context-shiftiness.

Since the structure of an argument manipulates the context in a way that determines the truth-conditional content of the component utterances, and does so as a matter of linguistic convention, we ought to reflect this effect in the logical form of an argument. The insight of the account proposed in Stojnić (2016c,a) and sketched above, is that the discourses don't merely express content, but also carry information about how to manipulate contextual parameters in a way that determines the content expressed. Following Stojnić (2016c,a), we can capture this by representing utterances as carrying two layers of content. On the one hand utterances express ordinary propositional content as in 3.1; on the other, they are represented as relations between contexts, conceived as our conversational scoreboards, capturing potential interpretive dependencies that an utterance brings about for the interpretation of subsequent context-sensitive expressions. Abstracting from the aspects of context that do not affect the interpretation of non-modal vocabulary, we think of context as a ranking of possibilities according to their relative prominence.

Formally, following Stojnić (2016c), we model a context as a set of assignment functions, where an assignment function is understood as a stack, specifying an ordered sequence of worlds.[21] So, a context is just a set of ordered sequences of worlds. We can think of it, intuitively as a ranking of propositions—sets of worlds—according to relative prominence.

Every stack in a context G begins with the 0^{th} position, as the top-ranked position on the stack. Then, for each position n, the position $n+1$ is one position lower in the ranking, representing diminishing prominence. Let 'G_i' denote the set of worlds that collects the worlds at the i^{th}

[21] Formally, a stack is an assignment function, mapping a finite convex subset of $\mathbb{N}$ to a set of worlds together with an undefined value $\perp$.

position of every stack in G.[22] Then, G_0 is the top-ranked proposition in G, and for each n, G_{n+1} is a proposition one step lower in the ranking. This allows us to keep track of the relative prominence of possibilities for modal anaphora.

We understand utterances, and the relations organizing them into a discourse, as carrying a dynamic layer of meaning. We capture this by specifying a dynamic language that models explicitly how the context evolves as the discourse unfolds, and then offering a translation of the relevant fragment of English into this language.[23] This will allow us to capture the dynamic aspect of meaning. Let our dynamic language have atomic expressions (propositional constants (p, q, r), and variables ($\mathbf{w}_i$ for $i \in \mathbb{N}$), conditions (propositional expressions comprising set of atoms closed under $\wedge$ and $\neg$), and update expressions, which we will define and describe shortly. For a given context G, let '$\mathbf{w}_i$' be a variable that stores a world at the i^{th} position of every stack in G. Then, relative to G, '$\mathbf{w}_i$' stores a set of worlds that comprises the worlds at the i^{th} position of every stack in G. This, in particular, will allow us to keep track of the propositions that have been introduced and promoted (as well as demoted), during the course of a discourse.

Our basic building block is the simplest update, which interprets the atomic formulae in our system, where we will interpret all the expressions that don't contain propositional subparts as atoms. The basic update just contributes the propositional content, but doesn't automatically interfere with the ranking. To capture this, we will reserve one designated position in the context, denoted by '*comp*', as a designated memory position, that stores the truth-condition of every clausal expression during the course of semantic composition.[24] The basic update then, just updates this position: it relates the input context G to an output context H, (relative to a world of evaluation w) just in case H differs from G in at most the value of '*comp*', and the value of '*comp*' is the proposition that 'p' expresses in G and w. That is, where G is the input context, H the output context, w the world of evaluation, and $G \underset{n}{\sim} H$ just in case G differs from H in (at most) the n^{th} position, we can define this update as follows:

Definition 5.1.

$[\![\langle comp := \phi \rangle]\!](w, G, H)$ iff $G \underset{comp}{\sim} H$ & $H_{comp} = [\![\phi]\!]^{G,w}$

Now we can use 5.1 to define more complex updates: the basic idea is that every non-atomic formula will be interpreted as making a contribution to the ranking of possibilities, and in the end storing the propositional content expressed by the overall utterance as a new value of '*comp*'. To see this, let's work through an update with a conditional.

Let '@E' denote the top-ranked epistemically live possibility (the top-ranked possibility that

[22] Thus, G_i is the proposition stored at the i^{th} position in G.

[23] See § A for an exposition of formal detailes.

[24] So, now, formally, a stack is just a function from a finite convex subset of $\mathbb{N}$ plus *comp* to a set of worlds together with an undefined value, $\perp$.

is a subset of epistemically accessible worlds) in a given context.[25] Then, a conditional utterance will be translated as an update with the following effect. Where G is an input context, the update first stores the proposition expressed by the antecedent (in G, relative to the world of evaluation w), as the value of '*comp*' and introduces the top-ranked epistemically live worlds in which the antecedent holds as the top-ranked possibility (pushing all other possibilities one position down). Relative to the thus obtained intermediate context (G''), it stores the proposition expressed by the consequent (in G'', relative to w), as the value of '*comp*' and introduces the top-ranked epistemically live worlds in which the consequent holds, pushing all other possibilities one position down in the ranking, resulting in the intermediate context (G''''). Lastly, the final output context (H) differs from the intermediate context G'''' only insofar as it stores the propositional contribution of the conditional as the value of '*comp*': as per 3.1, it stores the proposition true at a world w just in case all the epistemically accessible worlds from w in which the anaphorically retrieved restrictor and the antecedent hold are such that the consequent holds in them as well.[26] Thus, where K_1 and K_2 represent updates associated with the antecedent and the consequent, respectively, we define the update associated with the conditional as follows:

Definition 5.2.

$[\![\text{IF}(@E, K_1, K_2)]\!](w, G, H)$ iff there is a G', G'', G''' and G'''' such that
$[\![K_1]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $[\![K_2]\!](w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap [\![@E]\!]^{G'',w}$ & $G'''' \underset{comp}{\sim} H$ & $H_{comp} = Cond([\![@E]\!]^{G,w}, G'_{comp}, G'''_{comp})$

Since our examples involve negation and conjunction, we also need to introduce an update associated with negation and conjunction. '*not* ϕ' is interpreted as carrying the following truth condition: it is true (at a context and relative to a world w) just in case w is a non-ϕ world. The update associated with an utterance of '*not* ϕ' simply stores the complement of the truth-condition expressed by 'ϕ' in the input context (relative to w), as the value of '*comp*' of the output context. We define this update as follows:

Definition 5.3.

$[\![\text{NOT}(K)]\!](w, G, H)$ iff there is a G' such that $[\![K]\!](w, G, G')$ & $G' \underset{comp}{\sim} H$ & $H_{comp} = [\![\neg comp]\!]^{G',w}$, where $[\![\neg comp]\!]^{G,w} = \mathcal{D}_\omega \setminus [\![comp]\!]^{G,w}$, where $\mathcal{D}_\omega$ is a set of possible worlds provided by the model.

Conjunction, in turn, proceeds as follows: where G is an input context, the update first stores the proposition expressed by the first conjunct (in G, relative to the world of evaluation w), as the value of '*comp*' and introduces the top-ranked epistemically live worlds in which the first conjunct

[25] Formally, we include in the basic vocabulary of our dynamic language unary predicates and a unary operator '@'. Where 'E' is a unary predicate and '@' a unary operator, '@E' is an atom. '@E' is interpreted as taking a property denoted by 'E' and delivering the top-ranked proposition satisfying it, denoted by '@E'.

[26] As before, I abstract away from the contribution of the ordering source.

holds as the top-ranked possibility (pushing all other possibilities one position down). Relative to thus obtained intermediate context (G''), it stores the proposition expressed by the second conjunct (in G'', relative to w), as the value of '*comp*' and introduces the top-ranked epistemically live worlds in which the second conjunct holds pushing all other possibilities one position down in the ranking, resulting in the intermediate context (G''''). Lastly, the final output context (H) differs from the intermediate context G'''' only insofar as it stores the propositional contribution of the conjunction as the value of '*comp*': the proposition comprising the worlds in which both conjuncts hold. Thus, where K_1 and K_2 represent updates associated with the first and the second conjunct, respectively, we define the update associated with the conditional as follows:

Definition 5.4.

$[\![\text{AND}(K_1, K_2)]\!](w, G, H)$ iff there is a G', G'', G''' and G'''' such that $[\![K_1]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $[\![K_2]\!](w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap [\![@E]\!]^{G'',w}$ & $G'''' \underset{comp_0}{\sim} H$ & $H_{comp} = G'_{comp} \cap G''_{comp}$

Notice that our memory position, *comp* will always store the truth-conditional content expressed by an utterance in the final output context. For instance, the output of an update with a conditional will have the proposition expressed by the conditional as the value of this position. And similarly, negation stores its propositional contribution as the value of '*comp*' in the output context. Asserting this content, in turn, requires that the proposition expressed holds at the world of evaluation. In addition, it also makes the possibility associated with the asserted content prominent. We can capture this by ensuring that an assertion promotes the set of top-ranked epistemically live worlds in which the asserted content holds as a novel top-ranked possibility, and requires that the actual world be within that set:

Definition 5.5.

$[\![\text{ASSERT}(K)]\!](w, G, H)$ iff there is a G' such that $[\![K]\!](w, G, G')$ & $G' \underset{0}{\approx} H$ & $H_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $w \in H_0$.

Finally, we just need to add coherence. Elaboration has a two-fold contribution. First, it promotes the possibility elaborated on (demoting others, as always) and then it requires that the proposition expressed by the utterance that elaborates on it stands in Elaboration relation to it.[27] Putting all this together, where ϕ is a possibility, and K an update representing the utterance elaborating on ϕ, we can characterize the update associated with Elaboration as follows:

Definition 5.6.

[27] We can provisionally characterize Elaboration relation between propositions ϕ and ψ, $Elab(\phi, \psi)$, by requiring that it holds just in case ϕ and ψ are centered around the same event or entity, i.e. just in case the event or scenario described by ψ is a part of the event or scenario described by ϕ. (Cf. Hobbs (1979), Asher and Lascarides (2003).) The only thing that matters is the way in which the relation affects prominence. Ditto for other coherence relations.

$[\![\textsc{Elab}(\phi, K)]\!](w, G, H)$ iff there are G' and G'' such that $G \underset{0}{\approx} G'$ & $G'_0 = [\![\phi]\!]^{G,w}$ & $[\![K]\!](w, G', G'')$ & $G'' \underset{0}{\approx} H$ & $H_0 = G''_{comp}$ & $Elab([\![\phi]\!]^{G,w}, H_0)$.

Finally, we have to add Conclusion relation to this system, which passes down the body of information incorporating the information expressed by the premises (and requires that the premises and conclusion stand in Conclusion relation):[28]

Definition 5.8.

$[\![\textsc{Conclusion}(K_1, K_2)]\!](w, G, H)$ iff there is a G' such that $[\![K_1]\!](w, G, G')$ & $[\![K_1]\!](w, G', H)$ & $Conclusion(G'_0, H_0)$.

With this in place, let us work thorough the example in (3). Where r stands for "a Republican will win" and n for "Reagan will win", and a for "Anderson will win", the pattern will be associated with the following logical form:

(23) **Conclusion**(
Assert(**If**(@E, $\langle comp := r\rangle$, **Elab**(w_0, **If**(@E, **Not**($\langle comp := n\rangle$), **Elab**(w_0, $\langle comp := a\rangle$))))));
Assert(**Elab**(w_0, $\langle comp := r\rangle$)),
Assert(**If**(@E, **Not**($\langle comp := n\rangle$), **Elab**(w_0, $\langle comp := a\rangle$))))
)

The following are the key steps in (23). By 5.2, the first conditional update introduces the possibility corresponding to the set of top-ranked epistemically accessible worlds in the initial input context, in which the proposition expressed by the antecedent holds, i.e. the set of epistemically live worlds in which a Republican will win. The consequent provides an elaboration of this possibility; as a result, this possibility is promoted to prominence (as per 5.6). Furthermore, it is required that the possibility introduced by the consequent stands in the Elaboration relation to the possibility introduced by the antecedent, which at this point is the possibility ranked at the position 0 (and, so, denoted by the first occurrence of '$\mathbf{w}_0$' in (23)). Since the consequent contains a conditional, by 5.2 again, the antecedent of this embedded conditional introduces a possibility in which the winner is not Reagan relative to the current top-ranked possibility, the one promoted by Elaboration, i.e., relative to the possibility in which a Republican wins. The consequent of the embedded conditional further elaborates on this possibility, requiring that it be one in which Anderson is the winner. By 5.2 again, the whole conditional premise expresses the proposition that for all that is known, if the winner is a Republican and is not Reagan, the winer is Anderson. By 5.5, the assertion update

[28] Again, we don't need to worry about the details of the specification of the relational aspect of the Conclusion relation. Provisionally we can characterize it as follows:

Definition 5.7. $Conclusion(\phi, \psi)$ iff $[\![\phi]\!]^{w,G} \subseteq [\![\psi]\!]^{w,G}$ for all G, w.

requires that the conditional holds of the actual world and promotes the set of epistemically live worlds in which it holds. The second premise elaborates on top-ranked possibility in the context output by the conditional premise, that is, on the set of epistemically accessible worlds in which the conditional premise holds. This possibility is made prominent, and is further required that in it, the winner is a Republican. The assertion of the small premise requires that it holds of the actual world, and promotes the set of epistemically live worlds in which it holds. The conclusion relates to the premises by Conclusion relation, promoting the top-ranked possibility in which the premises hold, i.e. the set of epistemically accessible worlds in which a Republican wins. The antecedent of the conditional relative to this possibility introduces a possibility in which the winner is not Reagan. The consequent further elaborates on this possibility, requiring that it be one in which Anderson is the winner. By 5.2 again, the conclusion expresses the proposition that for all that is known, if the winner is a Republican and is not Reagan, the winer is Anderson. This is exactly what we expect.

We can also represent (4). As per our sketch in the previous section, where b stands for "the butler is innocent", g stand for "the gardener is innocent" and c for "the cook is guilty", (4-a) and (4-b) are represented, respectively, as (24) and (25):

(24) $\mathbf{Assert}(\mathbf{If}(@E, \mathbf{And}(\langle comp := b\rangle, \mathbf{Elab}(w_0, \langle comp := g\rangle), \mathbf{Elab}(w_0, \langle comp := c\rangle)))$

(25) $\mathbf{Assert}(\mathbf{If}(@E, \langle comp := b\rangle, \mathbf{Elab}(w_0, \mathbf{If}(@E, \langle comp := g\rangle, \mathbf{Elab}(w_0, \langle comp := c\rangle))))$

The key steps are as follows. In (24) the conditional introduces, relative to a top-ranked set of epistemically live worlds in the initial context a possibility in which the butler is innocent. The second conjunct in the antecedent further elaborates on this possibility, requiring that in it the gardener is innocent, as well. The consequent further elaborates on this possibility, and requires that it be one in which the cook is guilty. (24) is true just in case all the worlds within the set of top-ranked epistemically accessible ones, in which the butler and the gardener are innocent, the cook is guilty. The assertion update requires that this possibility holds in the actual world. (25), similarly introduces a possibility in which the butler is innocent, relative to the set of top-ranked epistemically accessible worlds in the input context. The consequent elaborates on this possibility. The Elaboration makes it prominent, and the conditional in the consequent is restricted by it. Relative to it, it introduces the possibility in which the gardener is innocent. The consequent of the embedded conditional further elaborates on it, and requires that it be one in which the cook is guilty. (25) is true just in case all the worlds within the set of top-ranked epistemically accessible ones, in which the butler and the gardener are innocent, the cook is guilty. The assertion update requires that this possibility holds in the actual world. It's easy to see that relative to the same input context (24) and (25) express the same propositional content.

What we have now is a way of representing a dynamic layer of content that encodes the effects of

discourse structure on the interpretation of modals and conditionals. These meanings can be seen as instructions to build propositional contents by manipulating parameters of the context. We also see that given this account of meaning, we can explain why *modus ponens* is valid: the dynamic meaning of the discourse—its logical form—determines propositional content of its component expressions in such a way that whenever the premises are true, the conclusion is true, as well. In a similar fashion, the dynamic meaning of (4-a) and (4-b) guarantee that the two discourses determine the same propositional content relative to any input context.

Now, as the context changes throughout the interpretation of an argument, our key question was, how to capture this kind of effect of a changing context while disallowing illicit shifts that lead to equivocation. Given our dynamic system, we can now answer this question. While before we said that the set of premises entails the conclusion just in case whenever the premises are true in a context, the conclusion is true in that context, we now say that the set of premises entails the conclusion just in case for any input context G, if updating with the premises leads to a non-empty output context, then updating that output context with the conclusion leads to a non-empty context. In other words, whenever the premises describe a possible update to a non empty context, updating that context with the conclusion leads to a non-empty context.

The two notions of validity correspond to each other: whenever the propositions expressed by the premises entail the one expressed by the conclusion, the argument will be translated by a valid formula in our dynamic language.[29]

6 The Argument Structure and Form

We have seen a systematic way of capturing the evolution of a context in a way that allows for a well-behaved account of argument validity. Unlike the traditional account, this account doesn't require a uniform antecedently fixed context that is held constant throughout the interpretation of an argument. This is a welcome result, because the account explains why we rule out the logical forms we have seen in (21) or (22). The account rather specifies a logical form of an argument that makes contextual dependences explicit, and consequently, predicts a certain underlying propositional pattern. It is in virtue of the form (23) has that it will be guaranteed to fit the propositional pattern $p \rightarrow q, p \models q$. That is, it is in virtue of the structure the argument in (23) has that it will express a propositional pattern that fits the logical form of *modus ponens*.

And the same point applies to (4): it is in virtue of the structure that renders the dynamic effects manifest that (4-a) and (4-b) express the same propositional content. And making this structure explicit explains why (4) cannot be associated with the—invalid—(22) pattern. It is, again, in

[29]The correspondence theorem is proven in Stojnić (2016c). In particular, Stojnić (2016c) proves that for a translation of a fragment of English into our dynamic system, for every way of constructing a dynamic formula there is a corresponding classical translation. The dynamic language preserves classical modal logic.

virtue of the structure the argument in (4) has that it will express a propositional pattern of the form $(p \wedge q) \rightarrow r \models p \rightarrow ((p \wedge q) \rightarrow r)$. Note that, while the pattern is valid, it is not the more general exportation pattern, namely $(p \wedge q) \rightarrow r \models p \rightarrow (q \rightarrow r)$.

So, on the resulting picture, arguments are individuated and get their meanings in part through their structure; these structures can be seen as recipes for building propositions, by manipulating context. Notice that the effect of the structure is indispensable. If we have considered (3) as a sum-total of sentences in no particular order, there is no guarantee that (3-a)–(3-c) would express a valid logical form. It is only when considered as a pointed piece of argument, with its component sentences related in a particular way, that we can guarantee that the argument in (3) is indeed a *modus ponens* argument. This means that individuation of arguments requires individuation of *structure* over and above the superficial pattern exhibited by the collection of sentences that comprise them.

And similarly for (4): it is only taking the specific underlying structure of (4-a) and (4-b) that the two will be equivalent. This is the sense in which our system doesn't validate the unrestricted pattern of exportation: due to the way the argument is structured, it corresponds to the logical form of $(p \wedge q) \rightarrow r \models p \rightarrow ((p \wedge q) \rightarrow r)$, not $(p \wedge q) \rightarrow r \models p \rightarrow (q \rightarrow r)$.

Now, one might complain at this point that we have taken the idea of context change too far. There are well known notions of informational consequence in the literature that propose to account for the data in a way that captures aspects of context-change, without relying on the notion of structure of the argument we have appealed to. So, one type of reaction to (3) and (4) is that one ought to take into account not only whether the propositions expressed by the premises are true at the world of utterance, but also whether the premises have been *accepted* in the context of utterance, where the context is understood as a body of information, determined by Stalnakerian common ground. Once the premises are accepted the context supports the conclusion as well. In assessing the conclusion, one is not allowed to suddenly disregard the premises that have already been accepted. Then we might understand the validity as requiring that the context is fixed in this further sense: we can't disregard the information that has, up to this point, been accepted in the context.[30]

This idea of validity, which tracks not the truth preservation in context, but rather the dynamic flow of information across an evolving context, has been developed in a framework of dynamic update semantics (Veltman, 1985, 1996; von Fintel, 2001; Gillies, 2010,?; Yalcin, 2007, 2012). The idea can be made precise as follows. Instead of understanding the meaning of a sentence as expressing a proposition relative to a context, and define validity as truth-preservation relative to a fixed context, we instead understand meaning as a context-change potential, a function that takes a context and delivers a context, incorporating the change the assertion of an utterance brings about to the

[30] Unless we signal in a special way that some information has been suspended, e.g., by exploiting counterfactual morphology.

context. To illustrate, suppose we start with a context, comprising a set of worlds, and assert 'John is happy' against that context. The result of the assertion will be to eliminate all the worlds in which John is not happy from that context. We can then think of the meaning of the utterance as an update function that takes the context, and delivers an output context as modified by the update associated with the utterance. Assuming a simple propositional language closed under conjunction and negation, we recursively define truth-conditions as follows:

Definition 6.1 (Basic Update Semantics)**.**

- $c[\![p]\!] = \{w \mid w \in c \cap p\}$[31]
- $c[\![\neg p]\!] = c \setminus c[\![p]\!]$
- $c[\![p \wedge q]\!] = c[\![p]\!][\![q]\!]$

Along with this definition of meaning goes a definition of truth: a sentence is true (accepted) in a context, just in case updating with that sentence does not change the context; intuitively, the context already supports the information expressed by the sentence. We say that ϕ entails ψ just in case updating with ϕ leads us to a context in which ψ holds. That is, a set of premises entails the conclusion just in case updating with the premises sequentially lands us in a context in which the conclusion is accepted.[32] A bit more formally:

Definition 6.2. Truth.

ϕ is true (supported) in c, $c \models \phi$, just in case $c[\![\phi]\!] = c$.

Then, we define entailment as follows:

Definition 6.3. Entailment.

$\phi_1, ..., \phi_n \models \psi$, just in case, for any c, $c[\![\phi_1]\!]...[\![\phi_n]\!] \models [\![\psi]\!]$.

Intuitively, an argument is valid just in case any context that accepts the information encoded in the premises, also accepts that of the conclusion.

Notice that even given the standard semantics for conditionals in § 2, this captures the explanation that is underlying (3): a context that contains only the worlds in which a Republican wins, and ones in which (3-a) is also accepted will all be ones in which (3-c) will be accepted as well. However, this doesn't yet predict why (4) is seamless. Given the standard semantics, a context that accepts (4-a) need not at the same time accept (4-b). To this end, the proposed account is

[31] I adopt the standard notation and write '$c[\![p]\!]$' where c is the context which is the argument of the update function $[\![p]\!]$. So, $c[\![p]\!] = \{w \mid w \in c \cap p\}$.

[32] A family of related dynamic notions of entailment is available. We focus on update-to-test entailment (instead of, e.g., test-to-test entailment, as we are interested in both the agglomeration and preservation of information in a context.

coupled with a specific understanding of the meaning of a conditional. To see what is at stake, ask what is the effect of an update with an indicative conditional. The answer proposed is that a conditional doesn't add some proposition to the common ground. Rather, intuitively, it tests whether the consequent holds in a context in which the antecedent has been hypothetically entertained. So, we interpret a conditional as an update that functions as a kind of a test on contexts: it first hypothetically updates the context with the antecedent, and then checks whether the consequent holds throughout such a hypothetically updated contexts. If it does, the update returns the original contexts; otherwise, the test fails, and we get an absurd state as an output. So formally, we add the following update to the basic update system in 6.1:

Definition 6.4. Codnitional Update.

$c[\![if(p)(q)]\!] = \{w \in c : (c[\![p]\!])[\![q]\!] = c[\![p]\!]\}$

A similar test-like interpretation is assigned to modal expressions as well. A modal utterance tests the context for whether it bares a particular relation to the prejacent specified by the force of the modal. For instance, an utterance of 'might p' tests the context to check whether it is compatible with p. If it does, the update returns the original contexts; otherwise, the test fails, and we get an absurd state as an output.

Given the above definitions, we can easily see that the 6.1 system, together with 6.4 will validate *mouds ponens*. In particular, in (3), whenever a body of information accepts the premises, it will accept the conclusion as well. The sense of infelicity in (3) is explained as above: the conclusion seems unsupported only if we subsequently revise our acceptance of the second premise. The proposed account validates import-export pattern as well, and in particular, it validates the problematic exportation direction: any context that accepts the conditional in (4-a) will also accept the one in (4-b).[33]

[33]See Gillies (2004). Suppose otherwise, that is, that for some context c the context accepts (4-a), but not (4-b). Then it follows that all the worlds in c in which neither the butler nor the gardener is the culprit, the cook is. Suppose there is at least some such world in c. Now, (4-b) first hypothetically adds the information that the butler is not the culprit to c, and then requires that the consequent—if the gardener didn't do it, the cook did—holds throughout the so updated c. Since we assumed that (4-b) is false, there has to be at least one world in thus hypothetically updated c, in which the consequent conditional is false. This means that once we further hypothetically update c with the information that the gardener is not the culprit, within thus narrowed down c there is at least one world in which the the cook is not the culprit. But, since (4-a) is true, all the worlds in thus narrowed down c—which are all the worlds in the original c in which the butler and the gardener are not culprits—are worlds in which the cook is the culprit. Thus, we get a contradiction.

On the other hand, suppose that there is no world in c in which neither the butler nor the gardener are the culprits. Then (4-a) is trivially true. But, then suppose that (4-b) is false. Then, updating the context c hypothetically with the information that the butler is not the culprit has to lead us to a non-empty context (as otherwise the conditional would be trivially true). This leaves us with the hypothetically narrowed down c in which the butler is not the culprit. Updating this context further with the information that the gardener is not the culprit also has to lead to a non-empty context, as otherwise the consequent would trivially hold, and the whole conditional would be true. But, this means that there has to be some worlds in w in which neither the butler nor the gardener are the culprits. This again leads to a contradiction. Hence, for all contexts, whenever the contexts accepts (4-a) it also accepts (4-b); the pattern is validated.

So, the dynamic account of logical consequence just sketched fares fairly well with respect to the data we observe. It does so by design: the notion of consequence is designed to preserve information agglomeration, or epistemic commitment. The premises entail the conclusion just in case for any body of information, if it is committed to the premises, it is committed to the conclusion as well.

Notice, however, that since this account doesn't attribute specific (nontrivial) propositional content to a conditional or a modal utterance, it doesn't adequately capture when an argument is employing an 'illicit' context shift of the sort that we see in (2), resulting in equivocation. This means that it might not allow us to individuate the logical forms of arguments finely enough to separate instances of equivocation from valid patterns. To see what is at stake, consider the following argument (Yalcin, 2007):

(26) There is an urn with a 100 marbles. 50 are small and blue, 30 are big and red, 10 are small and blue, and 10 are small and red.

a. If the marble is big, it is likely red.

b. It is not likely red.

c. So, the marble is not big.

This argument is problematic: clearly the conclusion doesn't follow. This is, of course, surprising on the standard account of conditionals. The standard account of conditionals validates *modus tollens*, so one cannot explain why is it that (26) is a bad argument, even if we commit to holding the context fixed: if all the closest worlds in which the marble is big are such that, relative to them, the marble is likely red, then, if the marble is *not* likely red (relative to the actual world), then the marble is big (relative to the actual world). But, this predicts that (26-a)—(26-c) should be a good inference. Yet, clearly it is not.[34]

However, if we endorse a dynamic account of a conditional and consequence in 6.4, then the view predicts that the argument is indeed invalid: if the subset of the context within which the marble is big, supports the update with the consequent (so that, given that subset, the marble is likely red), and the context as a whole supports the update with the claim that the marble is likely red, it doesn't follow that the context as a whole accepts that the marble is big. Now, since the inference pattern is bad, this might be considered a good result. However, as argued by Stojnić (2016c), if we understand the interpretation of modals and conditionals as sensitive to discourse

[34]Note that the probability modal 'likely' isn't crucial for the example:

(i) Whenever John is in France, he only ever visits Paris. But he rarely goes to France, and is almost always in Italy instead. We know he's in Europe now.

a. If John is in France he must be in Paris.
b. But it's not the case that he must be in Paris. (But, he might not be in Paris.)
c. So, it's not the case that he is in France.

structure, it becomes clear that the above inference doesn't exhibit the pattern of *modus tollens* to begin with.[35]

This is easy to see given our account above. The conditional premise introduces a hypothetical scenario, one in which the marble is big, against the prominent body of information—the epistemically live scenario including the background information in (26). The consequent of the conditional further elaborates on this possibility introduced by the antecedent, and Elaboration relation promotes the possibility elaborated on. The modal in the consequent is thus restricted by it, and hence receives a restricted reading: the marble is likely red given the epistemically live possibility in which it is big. The second premise stands in Contrast relation to the first one. Contrast requires that (26-a) and (26-b) provide contrasting bits of information about some body of information regarding the likelihood of the redness of the marble, given some or no assumption about its size. The body of information that the contrast is about is just the one relative to which the first conditional is interpreted, the one comprising epistemically live worlds in which the background information about the color and size distribution of the marbles holds. Contrast makes this body of information prominent, and the second sentence thus has to provide a contrasting bit of information about it regarding the likelihood of the marble being red. As a result, the modal in the small premise selects this body of information as its restrictor, and is thus interpreted as expressing the proposition that the marble is not likely red given this overall body of knowledge (given no particular assumption about its size).

This explanation predicts that, as a result of the import of discourse structure, the propositional content expressed by (26) doesn't fit the form of *modus tollens*: the small premise doesn't express a proposition negating the one expressed by the consequent of the big premise. The form expressed is $p \rightarrow q, \neg r \models \neg p$, a clearly invalid one. So, once we take into account the structure of the argument it is easy to see that in our system it will be invalidated. That however, doesn't threaten validity of *modus tollens*, as (26) is not an instance of *modus tollens* to begin with. An inference pattern that expresses the relevant propositional pattern, in turn, will be associated with a valid form.

While 6.1 treats modal and conditional sentences as sensitive to a background body of information they are interpreted against, it does not really treat them as context-sensitive in an important sense: while the truth of an update with a sentence varies depending on the input context, the account does not generally predict that the information carried by a sentence is affected by the potential interpretive dependences encoded in the input context (so, for instance, we do not capture the sense in which the change in context brought about by the antecedent changes the proposition expressed by the consequent). Yet this information is crucial for individuating the logical form of an argument: just as it is crucial for determining the information expressed by (2) that the premise and the conclusion in (2) do not express the same content due to the change in context that im-

[35] For the similar reason, the account in 6.1 doesn't distinguish the restricted pattern of exportation that (4) exhibits, from the unrestricted pattern, namely, $(p \wedge q) \rightarrow r \models p \rightarrow (q \rightarrow r)$.

pacts the interpretation of the pronoun 'he', so too, it is crucial to keep track of the difference in interpretation that affects the information expressed by the component expressions in (26), (3) and (4). And this in turn is crucial for the individuation of an argument: for exactly the same reason we say that (2) isn't a counterexample to the corresponding classical inference pattern, we shouldn't interpret (26) as an instance of modus tollens, or (4) as an instance of exportation.[36]

7 Conclusion

The traditional account of validity in a natural language requires that an argument is evaluated against a stable context which guarantees a non-equivocal interpretation of component utterances. This measure was supposed to prevent illicit context-shifting that can give rise to equivocation. This kind of ban on context-shifting was employed to explain why, e.g., (2) doesn't threaten the validity of the $p \models p$ inference.

But this universal ban on context-shifting is inadequate if the aim is to capture actual validity patterns in a natural language like English. We have seen that, often, it is precisely the change in context that underwrites a particular argument pattern: for instance, it is the change in context brought about by the small premise in (3) that is crucial for understanding (3) as a valid inference pattern.

Individuating argument forms not just as a set of premises and conclusion, but rather as partly individuated by their structure allowed for an account of validity that doesn't require a fixed context and uniform interpretation across an argument, yet allows us to distinguish shifts in contexts that the standard account would treat as 'illicit', underwriting the shifts in interpretation that undermine validity, as would a change in context in (2). The way we did this is by isolating a layer of meaning that governs the changes in context that affect interpretation. Utterances are interpreted as updates to the context, which update and potentially change interpretive dependences that transpire across the subsequent discourse.

While the examples we focused on involved modal language, the point generalizes to other cases as well. For instance, consider the following:

(27) John came in. He sat down. So, John came in and sat down.

Prima facie, (27) seems pretty good. But of course, it's not the case that for any model and *any* context in which the premises are true, the conclusion is, as well: just take a context in which

[36] Another way of looking at this is as follows: on my account, instead of thinking of updates as manipulating information, we are thinking of them as manipulating contextual parameters. In this way, they are building the contents expressed. And how they do so is sensitive not just to the contribution of individual utterances, but how they are structured into an argument. Thus, the argument structure in turn dictates whether given any input context, if the premises express a true proposition, the conclusion will, as well. As discussed above, whenever an argument expresses a sequence of updates that satisfies this requirement, it will be dynamically valid in our sense a well.

'he' is uttered while pointing at someone other than John, say Bill, and while Bill sat down, John remained standing.[37]

One way to handle this is to maintain that 'he' has a variable like behavior. So, in (27), it can be interpreted as co-indexed with the antecedent, 'John', thus receiving the dependent interpretation, picking out 'John'. But part of the problem is that a choice of the appropriate indices is standardly treated as underspecified by discourses like (27): it is precisely the context that is to tell us which index a pronoun is carrying.[38] So, the context then first has to do the work of fixing the underlying form of the argument. One could think of this role of context as that of *disambiguation*, where one form is selected out of the alternative possible ones determined by grammar. However, the pronoun itself doesn't seem to be ambiguous (after all, there is only one word 'he'), and it is unusual to treat examples like (27) as examples of structural ambiguity. So typically, it is assumed that grammar under-determines the form of utterances containing demonstrative expressions, and the context does the work of fixing the interpretation. On this account, the context isn't merely selecting one form among various complete ones, but rather plays a role in fixing a specific form to begin with. Only once the interpretation is fixed, the sequence can be evaluated for truth or validity.

But the account developed here allows us to maintain that the form of the argument in (27) depends on its underlying structure. If the second sentence follows the narrative about John, (27) will be associated with one form, the one which requires that the pronoun co-varies with 'John'. If, on the other hand, the second sentence is uttered in tandem with a pointing gesture, or perhaps contrastive stress, the form of the argument will be a different, invalid one.[39] The account thus provides us with resources to capture the effects of the changes in context that component expressions bring about, while at the same time allowing us to individuate argument forms finely enough that differences in the information expressed will transpire as differences in form.

[37] A similar point is illustrated by the following example:

(i) A man came in. He sat down. So, some man came in and sat down.

On a natural interpretation, the pronoun receives a 'discourse bound' reading, co-varying with the antecedent 'a man'. But if it is uttered with a pointing gesture towards a dog, Fido, it receives a deictic, referential interpretation.

[38] Note that things aren't as simple as differentiating between a bound and deictic use of a pronoun, and then maintaining that discourses like (27) are ambiguous between two different logical forms corresponding to two different assignment of indices to pronouns. For, 'he' can fail to be co-indexed with 'John' even when it's discourse bound by a different antecedent:

(i) Fido started pacing around 3pm. Suddenly, he ran to the door. John came in. He jumped at him.

Or similarly, consider the following:

(ii) A man met John. He greeted him. So, some man greeted John.

[39] This point is made formally precise in Stojnić, Stone, and Lepore (2017).

References

Asher, Nicholas and Lascarides, Alex. 2003. *Logics of Conversation*. Cambridge University Press.

Bittner, Maria. 2001. "Topical Referents for Individuals and Possibilities." In Rachel Hastings (et al.) (ed.), *Proceedings from SALT XI*, 36 – 55. CLC Publications, Cornell University, Ithaca.

—. 2011. "Time and Modality Without Tenses and Modals." In R Musan and M Rathert (eds.), *Tense Across Languages*, 147–88. Niemeyer, Tübingen.

—. 2014. *Temporality: Universals and Variation*. Wiley-Blackwell.

Brasoveanu, Adrian. 2010. "Decomposing Modal Quantification." *Journal of Semantics* 27:437–527.

von Fintel, Kai. 2001. "Counterfactuals in a Dynamic Context." In Michael Kenstowicz (ed.), *Ken Hale: A life in language*, 123–152. MIT Press.

Gibbard, Alan. 1981. "Two Recent Theories of Conditionals." In Robert Stalnaker William Harper and Glen Pearce (eds.), *Ifs*, 211–247. Reidel.

Gillies, Anthony S. 2004. "Epistemic Conditionals and Conditional Epistemics." *Noûs* 38:585–616.

—. 2010. "Iffiness." *Semantics and Pragmatics* 3:1–42.

Grice, H.P. 1989. "Indicative Conditionals." In *Studies in the Way of Words*, 58–85. Cambridge, MA: Harvard University Press.

Hobbs, Jerry R. 1979. "Coherence and Coreference." *Congitive Science* 3:67 – 90.

—. 1990. *Literature and Cognition*, volume (CSLI Lecture Notes 21). University of Chicago Press.

Kaplan, David. 1989a. "Aterthoughts." In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 565–614. Oxford University Press.

—. 1989b. "Demonstratives." In Almog Joseph, Perry John, and Wettstein Howard (eds.), *Themes From Kaplan*, 481–563. Oxford University Press.

Kehler, Andrew. 2002. *Coherence, Reference and the Theory of Grammar*. CSLI Publications.

Kehler, Andrew, Kertz, Laura, Rohde, Hanna, and Elman, Jeffrey. 2008. "Coherence and Coreference Revisited." *Journal of Semantics* 25:1–44.

Knott, Alistair. 1996. *A Data Driven Methodology for Motivating a Set of Coherence Relations*. Ph.D. thesis, Department of Artificial Intelligence, University of Edinburgh.

Kratzer, Angelika. 1983. "Conditionals." In Anne M. Farley, Peter Farley, and Karl-Erik McCullough (eds.), *Papers From the Parasession on Pragmatics and Grammatical Theory*, 1–15. Chicago: Chicago Linguistic Society.

Lepore, Ernest and Stone, Matthew. 2015. *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.

Lewis, David. 1975. "Adverbs of Quantification." In Edward L. Keenan (ed.), *Formal Semantics of Natural Language*, 178–188. Cambridge, MA: Harvard University Press.

Lewis, David K. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8:339 – 359.

Mann, William C and Thompson, Sandra A. 1988. "Rhetorical Structure Theory: Toward a functional theory of text organization." *Text* 8:243–281.

McGee, Van. 1985. "A Counterexample to Modus Ponens." *Journal of Philosophy* 92:462 – 471.

Roberts, Craige. 1989. "Modal Subordination and Pronominal Anaphora in Discourse." *Linguistics and Philosophy* 12:683 – 721.

Smyth, Ron. 1994. "Grammatical Determinants of Ambiguous Pronoun Resolution." *Journal of Psycholinguistic Research* 23:197–229.

Stojnić, Una. 2016a. "Content in a Dynamic Context." *Manuscript, Rutgers University* .

—. 2016b. *Context-Sensitivity in A Coherent Discourse*. Ph.D. thesis, Manuscript, Rutgers Universtiy.

—. 2016c. "One Man's *Modus Ponens...*: Modality, Coherence and Logic." *Philosophy and Phenomenological Research* doi:10.1111/phpr.12307.

Stojnić, Una, Stone, Matthew, and Lepore, Ernest. 2013. "Deixis (Even Without Pointing)." *Philosophical Perspectives* 27:502 – 525.

—. 2017. "Discourse and Logical Form: Pronouns, Attention and Coherence." *Linguistics and Philosophy* doi:10.1007/s10988-017-9207-x.

—. forthcoming. "Distinguishing Ambiguity from Underspecificity." In Ken Turner and Lawrence Horn (eds.), *An Atlas of Meaning*, volume Current Research in the Semantics/Pragmatics Interface. Leiden: Brill.

Stone, Matthew. 1997. "The Anaphoric Parallel Between Modality and Tense." *IRCS Report 97–06, University of Pennsylvania* .

—. 1999. "Reference to Possible Worlds." Technical Report 49, Rutgers University, New Brunswick, RuCCS.

Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.

—. 1996. "Defaults in Update Semantics." *Journal of Philosophical Logic* 25:225–261.

Wolf, Florian, Gibson, Edward, and Desmet, Timothy. 2004. "Discourse Coherence and Pronoun Resolution." *Language and Cognitive Processes* 19:665–675.

Yalcin, Seth. 2007. "Epistemic Modals." *Mind* 116:983 – 1026.

—. 2012. "A Counterexample to *Modus Tollens*." *Journal of Philosophical Logic* 41:1001 – 1024.

Zvolenszky, Zsófia. 2002. "Is a Possible-Worlds Semantics of Modality Possible? A Problem for Kratzer's Semantics." In Brandon Jackson (ed.), *Semantics and Linguistic Theory XII*, 339–358. Ithaca, New York: Cornell University Press.

A Appendix: Formal definitions[40]

In this appendix, I provide a dynamic language that formalizes my approach to modality with anaphora, and provide a translation from a fragment of English into this language.

A.1 Syntax:

In this section I specify the expressions of the language. We first start by listing the basic vocabulary:

- Propositional expressions: the elements of the set $\mathcal{C}$ of constants $(p, q, r...)$, and the elements of the set $\mathcal{V}$ of variables (*comp* and $\mathbf{w}_n$ for $n \in \mathbb{N}$).
- Unary predicates: P, Q, R
- Unary operator: @
- Update expressions: K, H
- Connectives: $\wedge, \neg$
- Identity: $=$

The following are atomic formulae (atoms) in our language:

[40] The dynamic language provided here is a slight extension of the on developed in Stojnić (2016c), expanded to capture relation of Conclusion.

- Propositional expressions are atoms.
- $@P$ is an atom, where P is an unary predicate.
- Nothing else is an atom in our language.

These are the conditions in our language:

- All atoms are conditions.
- $\phi = \psi$ is a condition, where ϕ, ψ are conditions. (Stands for identity.)
- $\neg\phi$, where ϕ is a condition. (Stands for negation.)
- $\phi \wedge \psi$ is a condition where ϕ, ψ are conditions. (Stands for conjunction.)

These are the update expressions:

- $\langle comp := \phi \rangle$ is an update expression, where ϕ is an atom.
- If ϕ is a condition, then $[\phi]$ is an update expression.
- $K; K'$ is an update expression, if K is an update expression and K' is an update expression.
- $\textsc{might}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\textsc{must}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\textsc{likely}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\textsc{if}(\phi, K_1, K_2)$ is an update expression, if ϕ is a condition and K_1 and K_2 are update expressions.
- $\textsc{and}(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.
- $\textsc{not}(K)$ is an update expression, if K an update expression.
- $\textsc{Assert}(K)$ is an update expression, if K is an update expression.
- $\textsc{Elab}(\phi, K)$ is an update expression, if ϕ is a condition and K an update expression.
- $\textsc{Conclusion}(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.
- $\textsc{Contrast}(K_1, K_2)$ is an update expression, if K_1 and K_2 are update expressions.

A.2 Models:

I define frames and models in the usual way:

- **A Frame** is a tuple $\mathcal{F} = \langle \mathcal{D}_w, \mathcal{D}_t = \{0,1\}, R, \mathcal{P} \rangle$ such that $\mathcal{D}_t$ is a domain of truth values ($\mathcal{D}_t = \{0,1\}$), $\mathcal{D}_w$ is a finite domain of possible worlds, $\mathcal{D}_t \cap \mathcal{D}_w = \emptyset$, with R, a (transitive and reflexive) accessibility relation defined over $\mathcal{D}_w$, and $\mathcal{P}$, a probability measure over $\mathcal{D}_w$, that maps each subset of $\mathcal{D}_w$ to $[0,1]$, satisfying the following constraints:

  i $\mathcal{P}(\mathcal{D}_w) = 1$

  ii $\mathcal{P}(p \cup q) = \mathcal{P}(p) + \mathcal{P}(q)$, when p and q are disjoint subsets of $\mathcal{D}_w$.

  iii $\mathcal{P}$ is a regular probability measure: if $p \neq \emptyset$ then $\mathcal{P}(p) > 0$.

- **A Model** is a pair $\mathcal{M} = \langle \mathcal{F}, \mathcal{I} \rangle$, where $\mathcal{F}$ is a frame and $\mathcal{I}$ an interpretation function, which assigns to each propositional constant p a subset of $\mathcal{D}_w$ and each predicate constant P a set of subsets of $\mathcal{D}_w$.

A.2.1 Truth-conditional contributions of modals and conditionals:

Let us define some meta-language abbreviations that will help us state the truth-conditions associated with updates associated with modals and conditionals. These correspond to propositions expressed by modals and conditionals.

- Where R is the accessibility relation, and $\mathcal{P}$ the probability measure over sets of possible worlds provided by the model:

Definition A.1. (Definition 3.1 in the text)

$Cond(p,q,r) := M(p \,\&\, q, r) = \{w \mid \forall w' : wRw'$, if $w' \in p \,\&\, w' \in q$, then $w' \in r\}$—*if* q, r, relative to p.

Definition A.2. (Definition 3.2 in the text)

$M(p,q) := \{w \mid \exists w' : wRw' \,\&w' \in p \,\&\, w' \in q\}$—*might* q, relative to some possibility, p.

Definition A.3.

$N(p,q) := \{w \mid \forall w' : wRw'$, if $w' \in p$ then $w' \in q\}$—*must* q, relative to some possibility p.

Definition A.4.

$P(p,q) := \{w \mid \mathcal{P}(\{w' \mid wRw' \,\&\, w' \in p \,\&\, w' \in q\}) \,/\, \mathcal{P}(\{w' \mid wRw' \,\&\, w' \in p\}) > .5\}$—*probably* q, given p.

A.3 Describing operations on stacks (sequences of worlds) and sets of stacks.

Here I define operations on stacks and sets of stacks, which I will use to define the semantics for our language later on. Formally, a stack is just a function from a finite convex subset of $\mathbb{N}$ plus *comp* to a set of worlds plus $\perp$, where '$\perp$' denotes an undefined value.[41] (I'll assume that '*comp*' is a designated position on the stack. Where the stack is intended to model prominence ranking, '*comp*' is not affecting the prominence ranking.)

- Where $m \in \mathbb{N}$, and i is a stack, i_m is the m^{th} member of the stack if m is within the domain of i, and $i_m = \perp$ otherwise. (i_{comp} is the member of the stack stored at the designated position *comp*.)— Selecting a member of the stack.
  - Where G is a set of stacks (i.e. a 'context'), g a stack, and u a world, $G_m = \bigcup_{g \in G}\{u | g_m \neq \perp \ \& \ g_m = u\}$, for $m \in \mathbb{N}$ or $m = comp$.—Getting the m^{th} element in the set of stacks G.
- For m, $n \in \mathbb{N}$, and a stack i, $i_{m,n}$ is a stack j defined on the set $\{0, ..., n-m\} \cup \{comp\}$ such that for $k \in \mathbb{N}, j_k = i_{(m+k)}$ if j is defined on k, and $j_{comp} = i_{comp}$.
  - Where G is a context, and g and j are stacks, $G_{m,n} = \bigcup_{g \in G}\{j | j = g_{m,n}\}$ and for $H = G_{m,n}, H_{comp} = G_{comp}$.
- For $m \in \mathbb{N}$, and a stack i, $i_m...$ is the stack j defined on the set $\{k \in \mathbb{N} \mid i \text{ is defined at} (m+k)\} \cup \{comp\}$ such that, for $k \in \mathbb{N}, j_k = i_{(m+k)}$ and $j_{comp} = i_{comp}$.
  - Where G is a context, and g, j are stacks, $G_{m...} = \bigcup_{g \in G}\{j | j = g_{m...}\}$ and for $H = G_{m...}, H_{comp} = G_{comp}$.
- If i is a stack with a finite domain with maximal element $k-1$ then for a stack j, $i + j$ is a stack h where, for $x \in \mathbb{N}, h_x = i_x$ if i is defined at x, and $h_x = j_{(x-k)}$ otherwise (and $h_{comp} = i_{comp}$).
- Where u is a world and i is a stack, u, i is a stack j, such that $j_0 = u$, and for all $n \in \mathbb{N}$, such that $n > 0, j_n = i_{(n-1)}$ if i is defined on n, and $j_n = \perp$ otherwise and $j_{comp} = i_{comp}$.— Appending to a stack.
  - Where G is a context, u is a world, and g, j are stacks, $G_{u...} = \bigcup_{g \in G}\{j | j = u, g\}$ and for $H = G_{u...}, H_{comp} = G_{comp}$.
- $g[n]g'$ iff $g_m = g'_m$ for $m \neq n$ (where $m, n \in \mathbb{N} \cup \{comp\}$).

[41] A set of numbers S is convex just in case if $x \in S$, $y \in S$ and $x < m < y$ then m is in S.

- $G \underset{n}{\sim} G'$ iff $\{g'|g[n]g', g \in G\} = \{g'|g[n]g', g \in G'\}$ (where $n \in \mathbb{N} \cup \{comp\}$).
- $G \underset{n}{\approx} G'$ iff $\{g_{0,n} + g_{n+1\ldots}|g \in G'\} = G$ and $G_{comp} = G'_{comp}$.

A.4 Semantics:

The Interpretation of Atoms: The interpretation of an expression e, relative to the interpretation function $\mathcal{I}$ a context G, and a world w:

- $[\![p]\!]^{G,w} = \mathcal{I}(p)$, if $p \in \mathcal{C}$.
  - Constants.
- $[\![\mathbf{w}_m]\!]^{G,w} = G_m$, if $\mathbf{w}_m \in \mathcal{V}$ and $m \in \mathbb{N}$
  - Variables.
- $[\![comp]\!]^{G,w} = G_{comp}$
  - A designated position on the stack.
- $[\![@P]\!]^{G,w} = \emptyset$ if $G_0 = \perp$, $[\![@P]\!]^{G,w} = G_0$, if $G_0 \in \mathcal{I}(P)$, and $[\![@P]\!]^{G,w} = [\![@P]\!]^{G_{1\ldots},w}$ otherwise.
  - Find the top ranked entity in G, satisfying P.

The Interpretation of Conditions:

- $[\![\phi = \psi]\!]^{G,w} = \mathcal{D}_\omega$, if $[\![\phi]\!]^{G,w} = [\![\psi]\!]^{G,w}$; $[\![\phi = \psi]\!]^{G,w} = \emptyset$, otherwise.
  - Identity.
- $[\![\neg\phi]\!]^{G,w} = \mathcal{D}_\omega \setminus [\![\phi]\!]^{G,w}$.
  - Negation.
- $[\![\phi \wedge \psi]\!]^{G,w} = [\![\phi]\!]^{G,w} \cap [\![\psi]\!]^{G,w}$.
  - Conjunction.

The Interpretation of Update Expressions

- $[\![\langle comp := p\rangle]\!](w, G, H)$ iff $G \underset{comp}{\sim} H$ & $H_{comp} = [\![p]\!]^{G,w}$
- $[\![[\phi]]\!](w, G, H)$ if and only if $H = G$ and $w \in [\![\phi]\!]^{G,w}$
- $[\![K; K']\!](w, G, H)$ iff $\exists G' : [\![K]\!](w, G, G')$ and $[\![K']\!](w, G', H)$

- The following are updates that describe how propositional content (A.2.1) in context is determined. Where p is a proposition (an anaphorically retrieved restrictor) and '@E' denotes the top-ranked proposition that is the subset of the epistemically accessible worlds:[42]

- $[\![\text{MIGHT}(\phi, K)]\!](w, G, H)$ iff there is a G' and G'' such that $[\![K]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $G'' \underset{comp}{\sim} H$ & $H_{comp} = M([\![\phi]\!]^{G,w}, G'_{comp})$

- $[\![\text{MUST}(\phi, K)]\!](w, G, H)$ iff there is a G' and G'' such that $[\![K]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $G'' \underset{comp}{\sim} H$ & $H_{comp} = N([\![\phi]\!]^{G,w}, G'_{comp})$

- $[\![\text{LIKELY}(\phi, K)]\!](w, G, H)$ if and only if there is a G' and G'' such that $[\![K]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $G'' \underset{comp}{\sim} H$ & $H_{comp} = P([\![\phi]\!]^{G,w}, G'_{comp})$

- $[\![\text{IF}(\phi, K_1, K_2)]\!](w, G, H)$ iff there is a G', G'', G''' and G'''' such that $[\![K_1]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $[\![K_2]\!](w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap [\![@E]\!]^{G'',w}$ & $G'''' \underset{comp}{\sim} H$ & $H_{comp} = Cond([\![\phi]\!]^{G,w}, G'_{comp}, G'''_{comp})$

- $[\![\text{AND}(K_1, K_2)]\!](w, G, H)$ iff there is a G', G'', G''' and G'''' such that $[\![K_1]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $[\![K_2]\!](w, G'', G''')$ & $G''' \underset{0}{\approx} G''''$ & $G''''_0 = G'''_{comp} \cap [\![@E]\!]^{G'',w}$ & $G'''' \underset{comp}{\sim} H$ & $H_{comp} = G'_{comp} \cap G''_{comp}$

- $[\![\text{NOT}(K)]\!](w, G, H)$ iff there is a G' such that $[\![K]\!](w, G, G')$ & $G' \underset{comp}{\sim} H$ & $H_{comp} = [\![\neg comp]\!]^{G',w}$

- $[\![\text{ASSERT}(K)]\!](w, G, H)$ iff there is a G' such that $[\![K]\!](w, G, G')$ & $G' \underset{0}{\approx} H$ & $H_0 = G'_{comp} \cap [\![@E]\!]^{G,w}$ & $w \in H_0$

In order to define the truth-conditions for updates associated with coherence relations, we assume the following abbreviations:

Definition A.5. $Elab(\phi, \psi)$ iff ϕ and ψ are centered around the same event or entity, i.e. iff the event or scenario described by ψ is a part of the event or scenario described by ϕ.

Definition A.6. $Conclusion(\phi, \psi)$ iff $[\![\phi]\!]^{w,G} \subseteq [\![\psi]\!]^{w,G}$ for all G, w.

Definition A.7. A formula, ϕ, is about of body of information θ iff, where G is the input context to ϕ, $\theta = [\![@E]\!]^{G,w}$, where '$E$' is a predicate denoting the property of being an epistemically accessible proposition, and thus, '@E' denotes the top-ranked epistemically accessible proposition. I use 'θ_ϕ' to denote the body of information that ϕ is about.

[42] For generality, I let the restrictor in the definition be any proposition p. However, as argued above, epistemic modals and conditionals select the top-ranked possibility in a given context ('@E') as their restrictor.

Definition A.8. *Contrast*(ϕ, ψ) iff ϕ and ψ describe contrasting information about some body of information regarding a common topic.

- $[\![\textsc{Elab}(\phi, K)]\!](w, G, H)$ iff there are G' and G'' such that $G \underset{0}{\approx} G'$ & $G'_0 = [\![\phi]\!]^{G,w}$ & $[\![K]\!](w, G', G'')$ & $G'' \underset{0}{\approx} H$ & $H_0 = G''_{comp}$ & $Elab([\![\phi]\!]^{G,w}, H_0)$.
- $[\![\textsc{Contrast}(K_1, K_2)]\!](w, G, H)$ iff there is a G' and G'' such that $[\![K_1]\!](w, G, G')$ & $G' \underset{0}{\approx} G''$ & $G''_0 = [\![\theta_{K_1}]\!]^{G,w}$ & $[\![K_2]\!](w, G'', H)$ & $[\![\theta_{K_1}]\!]^{G,w} = [\![\theta_{K_2}]\!]^{G'',w}$ & $Contrast(G''_{comp}, H_{comp})$
- $[\![\textsc{Conclusion}(K_1, K_2)]\!](w, G, H)$ iff there is a G' such that $[\![K_1]\!](w, G, G')$ & $[\![K_1]\!](w, G', H)$ & $Conclusion(G'_0, H_0)$.

A.5 Truth, validity, entailment.

- K is true, relative to a context G, a world w, and a model $\mathcal{M}$, if there is some H, s.t. $H \neq \emptyset$ and $[\![K]\!](w, G, H)$. K is false (relative to a context G, a world w, and a model $\mathcal{M}$,) otherwise.

- K is valid iff it's true in all models.
- K_1 entails K_2 iff for any model $\mathcal{M}$, any context G, and any world w if there is a G' such that $G' \neq \emptyset$ and $[\![K_1]\!](w, G, G')$, then there is a G'' such that $G'' \neq \emptyset$ and $[\![K_2]\!](w, G', G'')$.

A.6 Dynamic interpretation

I now give a dynamic translation for a fragment of English, specifying the updates associated with utterances containing modals and conditionals. I shall avail myself of abstract level of logical forms (ALFs) for the relevant fragment of English. The reader should bear in mind that we do not have to take a stand on the existence of a level of representation corresponding to ALFs. This level of representation is merely a dispensable convenience that helps specify the interpretation in a streamlined way.

A.6.1 Abstract Logical Forms (ALFs) for a Fragment of English:

Terms:

- Propositional constants from our base language in § A.1 (set $\mathcal{C}$).

Atoms:

- All terms are atoms, and nothing else is an atom.

ALFs:

- Atoms are ALFs.
- If ϕ and ψ are ALFs, then $might(\phi, \psi)$ is an ALF. (Stands for "it might be the case that ψ, given the restrictor ϕ".)
- If ϕ and ψ are ALFs, then $must(\phi, \psi)$ is an ALF. (Stands for "it must be the case that ψ, given the restrictor ϕ".)
- If ϕ and ψ are ALFs, then $likely(\phi, \psi)$ is an ALF. (Stands for "it's likely the case that ψ, given the restrictor ϕ".)
- If ϕ, ψ and γ are ALFs, then $if(\phi, \psi, \gamma)$ is an ALF. (Stands for "given the restrictor ϕ if ψ, then γ.")
- If ϕ and ψ are ALFs, then $and(\phi, \psi)$ is an ALF. (Stands for "ϕ and ψ.")
- If ϕ is an ALF, then $not(\phi)$ is an ALFs. (Stands for "Not ϕ.")
- If ϕ is an ALF, then $Assert(\phi)$ is an ALF. (Assertion operator—makes sure that the proposition ϕ is asserted.)

A.6.2 Dynamic Interpretation:

In this section, I provide a translation of the relevant fragment of English, into our dynamically interpreted language defined in § A.1–§ A.4. I'll assume the ALFs for the relevant fragment of English defined in A.6.1, (e.g. $might(\phi, \psi)$ for "it might be the case that ψ", where the modal is anaphorically dependent on $[\![\phi]\!]^{G,w}$, for an input context G.)

(Base case, where $T_d(\phi)$ is a translation of a formula ϕ into our dynamic system.)

- If ϕ is an atom, then $T_d(\phi) = \langle comp := \phi \rangle$.

(Recursive case)

- $T_d(might(\phi, \psi)) = \text{MIGHT}(\phi, T_d(\psi))$
- $T_d(must(\phi, \psi)) = \text{MUST}(\phi, T_d(\psi))$
- $T_d(likely(\phi, \psi)) = \text{LIKELY}(\phi, T_d(\psi))$
- $T_d(if(\phi, \psi, \gamma)) = \text{IF}(\phi, T_d(\psi), T_d(\gamma))$
- $T_d(and(\phi, \psi)) = \text{AND}(T_d(\phi), T_d(\psi))$
- $T_d(not(\phi)) = \text{NOT}(T_d(\phi))$
- $T_d(Assert(\phi)) = \text{ASSERT}(T_d(\phi))$