Conditional antecedents provide no evidence for a grammatical theory of scalar implicature

Daniel Lassiter
Stanford University

Draft of 11/19/13.
Comments welcome: danlassiter@stanford.edu.

1 Overview

Recent work in formal pragmatics has devoted much attention to the question of whether scalar implicatures are generated by conversational reasoning or by grammatical devices. The former approach has been dominant for some time, due in large measure to the work of Grice (1975, 1989). However, there has been a recent flourishing of grammatically-oriented alternatives in which scalar implicatures are generated by insertion of a silent exhaustification operator in the syntax. The most influential references are Chierchia 2006; Chierchia et al. 2012.

This paper discusses one of the most influential arguments that have been put forward in favor of a grammatical approach to scalar implicature, and shows that it a non-starter: data involving the behavior of scalar items in conditional antecedents fails to provide evidence of any kind in favor of the grammatical approach to scalar implicature. The arguments that have been put forward begin with the assumption that the examples in (1) should be infelicitous if the antecedents are not strengthened by (in effect) a silent version of *only*.

(1) a. If you ate some of the cookies and no one else ate any, there must be some left.
b. If it's warm, we'll go outside. But if it's hot we won't go outside.

Without silent exhaustification, it is claimed, (1a) (from Levinson 2000) is a contradiction and (1b) is an incoherent discourse. This does follow on the assumption, generally adopted by those who make this argument, that *if* denotes a material conditional; but this analysis has a number of widely-acknowledged problems and is no longer taken seriously in the literature on conditionals.

In fact, examples with this logical structure are widely discussed in the conditionals literature. A major desideratum for a theory of conditionals for some time has been to render coherent examples such as (2) — which mirrors, without a scalar item, the crucial entailment relationship between the antecedents in (1b).

(2) If this match is struck, it will light. But if it is soaked in water and struck, it will not light. (Goodman, 1947; Lewis, 1973)

Because of this and other well-known failures of monotonicity in conditional inferencing, virtually all current theories of conditionals are either semantically non-monotonic in the antecedent (Lewis, 1973, 1975; Kratzer, 1986), or contain pragmatic mechanisms which mimic non-monotonicity (von Fintel 1999; 2001). Adopting a better-motivated theory of conditionals thus leads us to expect that the examples in (1) should be fully acceptable under certain conditions. In this way, a widely-cited argument for the grammatical theory of scalar implicature is rendered inert.

An apparent problem for this reasoning is that the sentences in (1) are apparently equivalent to examples with an overt disjunction in the antecedent, which **are** contradictory:

(3) a. # If you ate some or all of the cookies and no one else ate any, then there must be some left.

b. If it's warm or hot, we'll go outside. # But if it's hot we won't go outside.

This is again a special case of a problem which is well-known in the literature on conditionals (e.g., McGee 1985; Edgington 1995). In order to explain similar inferences not involving scalar items, non-monotonic theories of conditionals are forced to posit additional mechanisms whereby conditionals with an overt disjunction in the antecedent (*If P or Q then R*) entail a conjunction of conditionals (*If P then R, and if Q then R*). §4 shows that such a mechanism is required regardless of one's position in the implicature debate, and that the existence of this mechanism explains the difference between (1) and (3) without recourse to silent exhaustification. Thus, despite appearances, examples like (3) do not tell us anything about which theory of scalar implicature is right.

In the concluding section I discuss the question of what type of arguments involving conditionals **could** be brought to bear on the scalar implicature debate. There is a small but potentially important difference between the theories that could be studied experimentally; however, there are also numerous potential confounds which are likely to make it extremely difficult to pry the theories' predictions apart. Only after careful consideration of these empirical and theoretical issues then will anyone be licensed to making strong theoretical claims on the basis of data involving scalar items in conditional antecedents.

2 The argument from conditional antecedents

On the Gricean picture, the inference from (4a) to (4b) is generated by reasoning about an interlocutor's beliefs and intentions. If a speaker utters (4a), we may be able to infer that (4b) is true using the following reasoning: if the speaker knew that (5) were true she would have said so, as long as certain enabling factors were present (e.g., a desire to be informative, lack of relevant politeness constraints, information and mental capacity to distinguish the relevant situations, an ability to remember and pronounce the word "hot", ...). Thus it stands to reason from her choice to utter (4a) that she thinks (5) is not true, and the conjunction of the literal content and the inferred negation of (5) is equivalent to (4b). However, if the relevant enabling factors are not present, then the inference to the negation of (5) is not appropriate, and the interpreter cannot go beyond the literal content of (4a).

(4) a. It's warm.
 b. It's warm but not hot.

(5) It's hot.

The grammatical approach to implicatures is different: a silent exhaustification operator *exh* can occur at any point in a syntactic tree. This operator "upper-bounds" the meaning of an expression, adding to it the negation of all of a set of alternatives ALT. (See Katzir 2007; Fox & Katzir 2011 for some ideas about how ALT is derived.) Restricting attention for simplicity to cases in which *exh* is applied to an expression of Boolean type (ending in t), we can define *exh* as in (6), where E is an expression whose denotation $\mathbf{E}$ is of type $\langle \alpha, t \rangle$.

(6) a. $[\![exh]\!] = \lambda B_{\langle \alpha,t \rangle} \lambda A_{\alpha}[B(A) \wedge \forall B'[B' \in ALT(B) \wedge B(A) \not\models B'(A) \rightarrow \neg B'(A)]$.
 b. $[\![[exh\ E]]\!] = \lambda A_{\alpha}[\mathbf{E}(A) \wedge \forall \mathbf{E}'[\mathbf{E}' \in ALT(\mathbf{E}) \wedge \mathbf{E}(A) \not\models \mathbf{E}'(A) \rightarrow \neg \mathbf{E}'(A)]$.

In the simplest case, if E denotes a proposition (type $\langle s, t \rangle$) then $[\![[exh\ E]]\!]$ picks out the proposition consisting of the conjunction of its denotation $\mathbf{E}$ and the negation of all alternatives which are not entailed by $\mathbf{E}$.

(7) u = "It's warm"; let $ALT = \{\mathbf{warm}, \mathbf{hot}\}$.
 a. Interpretation 1: [It [is warm]] $\Rightarrow$ **warm**
 b. Interpretation 2: [It [is [*exh* warm]]] $\Rightarrow$ $\mathbf{warm} \wedge \neg \mathbf{hot}$

On this theory the utterance "It's warm" is syntactically and semantically ambiguous, with reading (7a) leaving open the possibility that it is hot and (7b) excluding this possibility.

Both theories predict the same ambiguity for "It's warm", but they derive it differently. The predictions of the two accounts diverge when we consider cases in which an item such as *warm* appears in the scope of another operator — notably for us, in the antecedent of a conditional. The grammatical theory predicts two relevant readings (and many more irrelevant ones):

(8) If it's warm, we'll go outside.
 a. No *exh* in antecedent: roughly, "If it is not cool, we'll go outside"
 b. With *exh* in antecedent: "If it is warm but not hot, we'll go outside"

Absent additional mechanisms for deriving some kind of highly marked meaning shift (cf. Geurts 2010), a Gricean theory does not predict that (8b) should be a possible reading.

The crucial argument for the grammatical theory, originally noted by Levinson (2000), pp.206-7, is that reading (8b) is needed to account for cases such as (9), repeated from (1b). (This will be our running example throughout the paper.)

(9) If it's warm, we'll go outside. But if it's hot we won't go outside.

Levinson (2000) comments that "the conditional as a whole is assessed on the assumption that the *implicatures* (as well as the entailments) of the antecedent are satisfied by the model" (p.205). Sauerland (2010) likewise claims that this is one of (only) two "intuitively clear cases of embedded implicature", since the sentence would be a contradiction if the "not all" implicature were not incorporated as part of the truth-conditional content of the antecedent (pp.1-2). Chierchia et al. (2012) claim of a logically identical example that "[a] coherent interpretation, which is clearly possible, requires an embedded implicature". In his recent book, Chierchia (2013) uses analysis of a similar example involving disjunction[1] as a linchpin of his general proposal for the relation between syntax, semantics, and pragmatics (see also Chierchia et al. 2012, §2.2). According to Chierchia,

> [W]e *have* to ... add the implicature at the level of the embedded S, to weaken the interpretation. ... There is no other way to avoid the contradiction.

3 Scalar items in conditional antecedents

There is, of course, another way to avoid the contradiction: we can adopt a better analysis of conditionals. Neither Chierchia (2013) nor the other authors who have used conditional antecedents as an argument for the grammatical theory consider this possibility.[2] What is remarkable about this situation is that material conditional treatment of *if* is totally inadequate, and widely acknowledged to be so in the literature on conditionals. More to the point, the argument from conditional antecedents relies crucially on the assumption that conditionals are downward monotonic in the antecedent. However, almost all current theories of the semantics of conditionals are semantically non-monotonic, as defined and illustrated in (15).[3]

[1]Note that, unlike Chierchia et al. (2012); Chierchia (2013); Sauerland (2010), I have avoided using overt disjunctions in examples or paraphrases. This is because the use of an overt disjunction introduces additional pragmatic inferences which act as a confound: see §4 for discussion.

[2]The only exception that I am aware of is a brief, dismissive footnote in Chierchia et al. 2012. The footnote admits that adopting a non-monotonic theory of the conditional would invalidate their argument, but assert that the issue does not matter because the epistemic prediction that these theories make (discussed in some detail below) is "clearly" not correct. But this is not at all clear, and in any case the example that is discussed there has several confounds (see §5 below). Even if the empirical claim were compelling, the argument would not go through: as I discuss in §3.3, the only plausible monotonic theory available to us also fails to substantiate the argument that Chierchia et al. (2012) make based on examples like (1). This happens for principled reasons, because of the need to block inferences which are closely related to the ones in question but which do not involve scalar items. The argument thus depends for its teeth on assumptions about the semantics of conditionals which are demonstrably wrong.

[3]I focus in this paper on indicative conditionals, since they are the main source of examples in the existing literature on the grammatical theory of scalar implicature. The phenonema described seem to extend to other types of conditionals — not a surprising fact if basic logic of different conditionals types is the same, modulo tense/mood interactions. Given that the same word *if* is used in all types and the differences involve tense/mood/aspect and the presence or absense of *then*, this ought to be the default position, as e.g. Stalnaker (1975); Strawson (1986); Edgington (1995) discuss. As Strawson (1986) puts it, "[I]t is obvious that the least attractive thing one could say about the *difference* between the two remarks [a past-tense indicative

(10) Downward monotonic semantics for *if*: If P entails Q, then *If Q, R* entails *If P, R*.

a. *It is hot* entails *It is warm.*

b. So, *If it is warm, we will go out* entails *If it is hot, we will go out.*

(11) Non-monotonic semantics for *if*: there may be P and Q such that P entails Q, but *If Q, R* is true while *If P, R* is false.

a. *It is hot* entails *It is warm.*

b. Nevertheless, *If it is warm, we will go out* and *If it is hot, we will not go out* may be jointly satisfiable.

3.1 Downward monotonic theories

The material conditional analysis has it that *If P then Q* is equivalent to *Either not-P or Q*, i.e., *not both P and not-Q*. This theory is fairly easy to dispense with. As Égré & Cozic (to appear) put it in a recent survey, this treatment "both under- and over-generates with regard to very typical inferences that we make in natural language". For example, Grice (1989) argues that his theory of pragmatics could account for one counter-intuitive aspect of the material conditional analysis, its truth when the antecedent is false; however, he also noted that it makes contradictory predictions about certain conditionals containing probability operators (pp.78-9; see also Kratzer 1986). A related observation is due to Lewis (1975), who pointed out that the material conditional treatment delivers incorrect truth-conditions for sentences with adverbs of quantification like (12), whether the adverb is given scope *in situ* (12b) or scoped wide (12c).

(12) If it rains, I usually take my umbrella.

a. Correct interpretation: "In most situations which involve both me and rain, I take my umbrella."

b. MC-narrow: "Either it doesn't rain, or I usually take my umbrella."

c. MC-wide: "Usually, either it doesn't rain or I take my umbrella."

MC-narrow will be true if it does not rain, or if I usually take my umbrella — even if I *never* take my umbrella when it actually does rain. MC-wide is no more plausible: it would also be true if I never take my umbrella in rain-situations, as long as rain is uncommon. As Lewis observes, deriving the correct interpretation of (12) requires treating the conditional antecedent as restricting the domain of quantification for *usually*, as in paraphrase (12a).

More importantly for our purposes, there are powerful objections not only to the material conditional treatment of *if*, but to **any** downward monotonic semantics for *if*. Notable

conditional and a related counterfactual] is that ... the expression 'if ... then ... ' has a different meaning in one remark from the meaning which it has in the other." Surprisingly, this has not been considered obvious by many philosophical logic: for instance, compare the three rather different logics that Lewis offers for three types of counterfactuals, indicatives containing an adverb of quantification, and plain indicatives in Lewis 1973, 1975, 1976. The empirical and logical issues to be dealt with here would no doubt be rather more complex if the non-unified approach position were to turn out correct; but our main point would continue to hold at least for indicative conditionals.

problems include the fact that a downward monotonic semantics validates Strengthening the Antecedent (13) and Conditional Transitivity (15). As a result, a monotonic theory will validate the crazy inferences in (14) and (16) (unless additional pragmatic mechanisms are assumed, cf. §3.3 below).

(13) **Strengthening the Antecedent**: *If P then R* entails *If P and Q, then R* for arbitrary Q.

(14) Counter-examples to **Strengthening the Antecedent**:
   a. If this match is struck, it will light. So, if this match is soaked in water and struck, it will light. (Goodman, 1947; Lewis, 1973)
   b. If John stole the earrings, he must go to jail. So, if John stole the earrings and then shot himself, he must go to jail. (von Fintel 1999)

(15) **Conditional Transitivity**: *If P then Q* and *If Q then R* jointly entail *If P then R*.

(16) Counter-examples to **Conditional Transitivity**:
   a. If Bill gets the promotion, Mary will resign her job. If Mary dies before the decision is made, Bill will get the promotion. So, if Mary dies before the decision is made, Mary will resign her job. (mod. from Edgington 1995)
   b. If I quit my job, I won't be able to afford my apartment. If I win a million, I will quit my job. So, if I win a million, I won't be able to afford my apartment. (Kaufmann, 2005)

See e.g. Edgington 1995; Bennett 2003 for discussion of these and many more problems involving inference patterns incorrectly licensed by a downward monotonic semantics for conditionals.

3.2 Non-monotonic theories

Because of these and related data, "in the modern semantic and philosophical literature on conditionals, it is now taken for granted that conditionals are not monotonic in their antecedent" (von Fintel 1999, p.136). On these accounts, the bizarre inferences in (14) and (16) are not valid, and the examples with which we started in (1) (repeated here) are not semantic contradictions.

(17) a. If you ate some of the cookies and no one else ate any, then there must be some left. (Levinson, 2000)
   b. If it's warm, we'll go outside. But if it's hot we won't go outside.

There are numerous highly-developed theories falling into this category, any of which is capable of dealing with the data just considered. I'll give a brief and informal description of a few such theories, focusing on features that will be useful to us later, and with no pretense of completeness or absolute faithfulness to the original sources. The reader is strongly encouraged to consult the references cited for the full story.

Ordering-based theories (Stalnaker 1968, 1975; Schlenker 2004; and — for counterfactuals — Lewis 1973). The basic idea of theories in this class is that we begin with an ordering of possible worlds $\succcurlyeq_w$, assumed to be (at least) reflexive and transitive. $w_1 \succcurlyeq_w w_2$ indicates that w_1 is "at least as close" to world w as w_2 is. Stalnaker (1968) states the truth-conditions of a conditional in terms of this concept of closeness, informally glossed as follows (p.102):

> Consider a possible world in which A is true, and which otherwise differs minimally from the actual world. *"If A, then B" is true (false) just in case B is true (false) in that possible world.*

Stalnaker (1968) does not commit to an analysis of the concept of "closeness" in this paper, but later (in Stalnaker 1984) invites an epistemic interpretation, speaking of various uses of conditionals in terms of "projections of epistemic policy onto the world" (p.x). For simplicity, and because it seems to fit most uses of a conditional without an overt modal or adverb of quantification, let's fix the epistemic interpretation as our main object of interest.

Stalnaker also assumes that there is a unique "closest" world, but others have wished to relax this assumption. If so, we may substitute:

(18) *If P then Q* is true in w if and only if Q is true in all of the the P-worlds which are closest to w.[4]

This is cashed out as follows: a world w_1 is among the P-worlds closest to w just in case there is no P-world w_2 strictly closer to w than w_1 is $[\neg\exists w_2 : w_2 \succcurlyeq_w w_1 \wedge w_1 \not\succcurlyeq_w w_2]$. There will be multiple such worlds if there are distinct worlds w_3, w_4 which are equally distant from w $[w_3 \succcurlyeq_w w_4 \wedge w_4 \succcurlyeq_w w_3$, or if the order is not connected, so that there are multiple distinct sets of "closest" worlds in its various branches.

This is still quite informal, but it is enough for our purposes. The first question to ask is whether this semantics meets the minimal criteria which we established in the last section, and used to reject downward monotonic theories: does it invalidate Conditional Transitivity and Strengthening the Antecedent? As Figure 1 illustrates, it does. Conditional Transitivity fails because the closest Q-worlds may be more distant than the closest R-worlds; the set of closest worlds verifying the antecedent will thus differ when we evaluate the two conditionals. In this case, *If Q then P* may come out true if P holds in the (distant) Q-worlds, and *If P then R* may come out true if R holds in the (close) P-worlds; nevertheless, R may fail to hold in the (distant) Q-worlds. No contradiction arises.

Strengthening the Antecedent fails for similar reasons. R could easily hold in the nearest P-worlds, but fail to hold in the nearest $P \wedge Q$-worlds, as long as all of the closest P-worlds happen to be $P \wedge \neg Q$-worlds. In this case, *If P then R* is true but *If P and Q then R* is false on Stalnaker's semantics.

The latter observation brings us back to an issue which is crucial for the arguments involving scalar implicature. Any model which verifies both (a) *If P then R* and (b) *If P*

[4]Even more complicated paraphrases are necessary if we want to allow for infinite sequences of ever-closer P-worlds. It isn't necessary to worry about this complication here, since it won't affect the uses to which we put these theories.

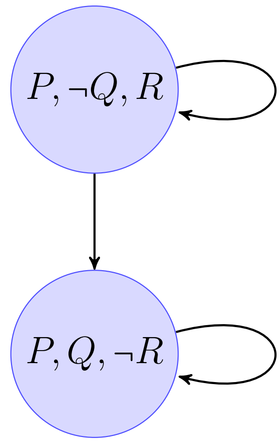

Figure 1: A tiny model exemplifying the failure of (a) Strengthening the Antecedent and (b) Conditional Transitivity in Stalnaker's semantics for conditionals. The world higher in the figure is "closer" to the evaluation world w in the sense discussed in the text. On Stalnaker's semantics, the model (a) verifies both *If P, R* and *If P and Q, not R*; (b) verifies *If P, R* and *If Q, P* but falsifies the transitive inference *If Q, R*.

and Q then not-R must have the property that the closest $P \wedge \neg Q$-worlds are closer than the closest $P \wedge Q$-worlds. For if not, it would not be the case that *all* of the closest P-worlds verify R: some of the closest P-worlds would be $P \wedge Q$-worlds, and we know from the truth of (b) that all of these closest such worlds make R false.

Continuing with the epistemic interpretation, we may wish to interpret "true in closer worlds" as "more likely" (cf. Kratzer 1991; but see also Lassiter 2010, 2011, 2014; Yalcin 2010 for objections to this interpretation). If so, we can extract from the ordering-based theory of conditionals a strong prediction: counter-examples to Strengthening the Antecedent will always have a particular epistemic profile.

(19) **Epistemic Inference**: *If P then R* and *If P and Q then not-R* jointly entail the following: *(P and not-Q) is more likely than (P and Q)*.

Given the failure of Strengthening the Antecedent in Stalnaker's semantics, it is no surprise that the examples from §1 involving scalar items in the antecedent of a conditional are no longer problematic. The logical relationship between *warm* and *hot*, or *some* and *all*, is precisely parallel to the one between P and $P \wedge Q$ (when P and Q are logically independent, that is). The two sentences in (17b) are jointly satisfiable on Stalnaker's semantics: all that is required is that the closest worlds in which it is warm-but-not-hot ($P \wedge \neg Q$) be closer than the closest worlds in which it is hot ($P \wedge Q$). And on an epistemic interpretation of closeness, we again derive a prediction: (17b) and similar examples should be acceptable only if it is more likely that the "strengthened" meaning of the weak scalar item is true than it is that the stronger scalar item is true. This prediction follows from the meaning of the conditional, without any special devices for manipulating the interpretation of scalar items.

The restrictor theory (Lewis, 1975; Kratzer, 1986; Egré & Cozic, 2011) A closely related theory is due to Kratzer (1986). Kratzer argues, generalizing a proposal by Lewis (1975), that we should think of conditionals as devices for restricting the domain of quantification of some operator. When there is an overt modal or adverb of quantification, the conditional (usually) restricts its domain; if not, it restricts a silent epistemic modal which is spelled out using Kratzer's (1981; 1991) theory of modality. Let f and g pick out the "modal base" and "ordering source", respectively; both are contextually provided functions from worlds to sets of propositions. Here are the truth-conditions for the case in which no overt operator is present:

(20) *If P then Q* is true in w, relative to f and g, just in case *must Q* is true in w relative to f_P^+ and g, where $f_P^+(w') = f(w') \cup P$ for any w'.

We have to add a condition if we want to ensure that the *must* invoked in the definition receives an epistemic interpretation; see Kratzer's papers for the details.

This theory of conditionals also fails to validate Conditional Transitivity and Strengthening the Antecedent, and it also validates the Epistemic Inference (19). To see why, let's examine a simplified version of Kratzer's theory of epistemic *must*.[5] The modal base f applies to the evaluation world w to produce a set of propositions representing the propositions known at w; we then take the conjunction of these propositions, $\bigcap f(w)$, to get a domain of "epistemically possible" worlds. *Must* is interpreted, as usual, as a universal quantifier; but it quantifies not over the entire set $\bigcap f(w)$ but over a distinguished subset: the worlds which satisfy as many as possible of the propositions in $g(w)$.[6] We use $g(w)$ to induce an order over worlds as follows:

(21) $w_1 \succcurlyeq_w w_2$ if and only if $\{p | p \in g(w) \wedge w_1 \in p\} \supseteq \{p | p \in g(w) \wedge w_2 \in p\}$.

That is, w_1 must satisfy all of the propositions in $g(w)$ that w_2 does, and possibly more, in order for $w_1 \succcurlyeq_w w_2$ to be true. Kratzer (1991) uses this order to give truth-conditions for *must* (presented here in the simplified form made possible by the limit assumption).

(22) *must P* is true at w iff all of the $\succcurlyeq_w$-undominated worlds are P-worlds. A world u is $\succcurlyeq_w$-undominated iff $u \in \bigcap f(w)$ and there is no $v \in \bigcap f(w)$ such that $v \succcurlyeq_w u$ but $u \not\succcurlyeq_w v$.

Conditional Transitivity and Strengthening the Antecedent fail in this semantics for essentially the same reason that they do in Stalnaker's theory. We'll focus on Strengthening the Antecedent, leaving the extension to Conditional Transitivity to the reader. The key observation is that *must* quantifies not over all worlds in the modal base, but over a (usually strict) subset of these worlds which are distinguished by being true in some maximal

[5] Again, I'm simplifying here by making the assumption that there is a unique set of "undominated"/"closest" worlds. As far as I can see this assumption does not affect the issues of interest involving implicature.

[6] That is, as long as $g(w)$ contains at least one proposition which is logically independent of $\bigcap f(w)$. With empty or otherwise trivial $g(w)$, Kratzer's *must* behaves like the $\Box$ of modal logic, with $\bigcap f(w)$ acting as $\{w' | wRw'\}$.

subset of the ordering source propositions. Suppose that *If P then R* and *If P and Q, then not-R* are both true at some world w. The first conditional is true just in case R is true in all of the worlds which are undominated relative to $\bigcap(f(w) \cup P) = \bigcap f(w) \cap P$. The second is true just in case R is true in all of the worlds that are undominated relative to $\bigcap(f(w) \cup (P \cap Q) = \bigcap f(w) \cap P \cap Q$. This is clearly consistent: all we need is an ordering source in which the highest-ranked P-worlds in $\bigcap f(w)$ satisfy R, while the highest-ranked $P \wedge Q$-worlds in $\bigcap f(w)$ satisfy $\neg R$. This will be the case, for example, whenever (a) Q and R are mutually exclusive, and (b) R is in the ordering source and $\neg R$ is not, or R is entailed by a (conjunction of) proposition(s) in the ordering source and $\neg R$ is not.

It follows immediately from this description that *If P then R* and *If P and Q, then not-R* can both be true only in models in which every $P \wedge Q$-world in $\bigcap f(w)$ is strictly dominated by some $P \wedge \neg Q$-world in $\bigcap f(w)$. Interestingly, this entails that *P is more likely than Q* comes out true on Kratzer's (1991) definition:

(23) a. *P is as likely as Q* is true at w iff, for every Q-world v in $\bigcap f(w)$, there is a P world u in $\bigcap f(w)$ such that $u \succcurlyeq_w v$.

b. *P is more likely than Q* is true at w iff both *P is as likely as Q* and *Q is not as likely as P* are true at w.

(23b) is true whenever all $P \wedge Q$-worlds in $\bigcap f(w)$ are weakly dominated by some $P \wedge \neg Q$-world in $\bigcap f(w)$, and at least one $P \wedge Q$-world is **strictly** dominated by a $P \wedge \neg Q$-world. This is entailed by the stronger condition that we derived above. Thus Kratzer's semantics for conditionals also makes the epistemic prediction: Strengthening the Antecedent is not valid, but its counter-models will always have a particular epistemic profile.

Most importantly for us, Kratzer's semantics agrees with Stalnaker's in its predictions about the special cases involving scalar items which proponents of the grammatical theory of scalar implicature have attempted to use as evidence for their position. Since *If it's warm we'll go out, but if it's hot we won't* is not a contradiction on Kratzer's theory, the fact that it is intuitively acceptable is not a problem if her theory is correct. And, again, the theory predicts unambiguously that this sentence can be true only if it's more likely that it will be warm-but-not-hot than it is that it will be hot.

Probabilistic theories: e.g. (Adams, 1975, 1998; Edgington, 1995; Kaufmann, 2005). A further important class of theories are those which draw a close connection between conditional probability and the semantics of conditionals. This is a vast and complex literature, complicated further by the fact that most theories of this type do not assign truth-conditions of the familiar type to conditional sentences. However, as Adams (1975, 1998); Edgington (1995) discuss in some detail, probabilistic theories do not validate Strengthening the Antecedent or Conditional Transitivity (in the extended sense of "validate" which they employ). Failing to validate Strengthening the Antecedent, they can also make sense of (17b) and the like without special mechanisms. Here again, arguments for the grammatical theory on the basis of data involving conditional antecedents thus fail if this is the right approach to conditionals. In the Appendix, I prove that the Epistemic Inference (19) accompanies counter-models to Strengthening the Antecedent in Edgington's theory as well, as long as the relevant conditionals have even moderately high probability.

3.3 Semantically monotonic, pragmatically non-monotonic (von Fintel 1999; 2001)

Von Fintel (1999; 2001) describes a theory which is semantically monotonic, in that the conditional is analyzed as a universal quantifier over the antecedent-satisfying worlds in a contextually given domain D_C. Since the universal quantifier is downward monotonic in its restriction, so is this theory of conditionals. This analysis is motivated (inter alia) by the desire to account for the presence of NPIs in the antecedents of conditionals; but it would seem to re-introduce the problems with Strengthening the Antecedent and Conditional Transitivity.

In order to avoid validating these inferences, von Fintel proposes that D_C is able to shift rapidly in the course of conversation, and in particular that it is always able to expand to include more remote possibilities, while contraction is difficult. So, for example, the first sentence in example (14a) (repeated here as (24)) is true on this account only if D_C excludes (as being remote or implausible, say) worlds in which the match is soaked in water or in some other special way disabled.

(24) If this match is struck, it will light. So, if this match is soaked in water and struck, it will light. (Goodman, 1947; Lewis, 1973)

However, the follow-up utterance triggers expansion of the domain D_C to D_C^+, which is required to include (at least) the highest-ranked worlds in which antecedent is true.[7] Clearly not all of the worlds in D_C^+ in which the match is soaked in water and struck will also be worlds in which the match lights. Thus, the second conditional is false relative to the expanded domain D_C^+, and so the first sentence is true while the second is false. This makes the reasoning in (14a) appear invalid, even though it would be valid if the domain were held constant. The theory is thus semantically monotonic but pragmatically non-monotonic.

If von Fintel's theory turns out to be correct, it is little help for proponents of the grammatical theory of scalar implicature. For (17b), we can tell a domain-shifting story which is precisely parallel to the account of (14a): in the initial context in which "If it's warm..." is uttered, worlds in which it is hot are excluded from D_C in the initial context (as being remote or implausible, say). But the follow-up "If it's hot, ..." triggers expansion of the domain to some D_C^+ which also includes worlds in which it's hot. Thus the reasoning is invalid despite the semantic monotonicity of the conditional.

The only potentially difficult case is (17a), and whether this is a problem depends on your analysis of *must*: if it quantifies over a subset of "best" worlds in the modal base (following in the essentials Kratzer 1981, 1991), then there is no contradiction here either. Again, we need only assume that the set of worlds quantified over does not contain any worlds in

[7] Von Fintel spells out "highest-ranked" using the modal base/ordering source apparatus of Kratzer (1981, 1991), described in the previous subsection.

which it is hot. This example would be problematic if both (a) *must* is a quantifier over all worlds in $\bigcap f(w)$ (as von Fintel & Gillies (2010) argue, but see **?**); and (b) the modal base were for some reason unable to shift mid-sentence. However, von Fintel (1999; 2001) argues that conditional domains of quantification are able to shift in this way, and I see no reason to doubt his logic in the related case of epistemic modals. The prospects for extracting a contradiction from examples involving scalar items in conditional antecedents thus remain extremely slim, even if von Fintel is correct about the semantic monotonicity of conditionals.

4 Overt disjunctions

Antecedent disjunctions (with or without scalar items) are problematic for all of the best theories of conditional semantics. Here is a typical example, taken from the web.

(25) I let him come to me only and fly to me which he does all the time. But if I or especially anyone else tries to pick him up, he pecks.

From (25) we infer categorically that, if the speaker (the bird's owner) picks the bird up, it pecks; and that if anyone else tries to pick up the bird in question, it pecks. As Edgington (1995) points out, this is surprising.[8] Of the theories we've considered, the ones that are able to deal with Strengthening the Antecedent and Conditional Transitivity uniformly assign this sentence a weaker truth-condition: if one disjunct happens to be true in closer worlds than the other, the consequent simply has to be true in the closest worlds satisfying that disjunct. On an epistemic interpretation of "closeness", this means that only the **more likely** of the disjuncts is relevant to evaluating the truth of the conditional sentence; the other is simply ignored. Thus, for (25), it follows from a plausible background assumption (the bird's owner picks up the bird more often than others do) that (25) tells us nothing about what happens when people other than the bird's owner try to pick it up. Instead, the sentence should be contextually equivalent to *If I try to pick him up, he pecks.*

Here is an invented example, which I'm imagining as directed by a biology professor to his fieldwork minions. (*Teinopalpus imperialis* is an extremely rare butterfly.)

(26) a. If you find a butterfly, you should take a note and tell me about it on Monday. But if you find a specimen of *Teinopalpus imperialis*, you should call me immediately.

b. If you find a butterfly or indeed a specimen of *Teinopalpus imperialis*, you should take a note and wait until Monday to tell me about it. # But if you find a specimen of *Teinopalpus imperialis*, you should call me immediately.

Otherwise successful theories of conditional semantics lead us to expect that (26a) and (26b) should be of equal status, and that both should be acceptable only if finding a specimen

[8] Edgington (1995) mentions this point as an objection to Stalnaker's theory, but it applies equally to Kratzer's restrictor theory and Edgington's own probabilistic theory. The application to Kratzer's theory is clear, due to its close structural resemblance to Stalnaker's. In Edgington's case, a rational person can have high credence in *If P or Q then R* and have low credence in *If Q then R*, because $p(R|P \vee Q)$ can be high when $p(R|Q)$ is low. This will be the case, for example, if $p(R|P)$ is high and $p(Q|P \vee Q)$ is low.

of *Teinopalpus imperialis* is less likely than finding some other kind of butterfly. These predictions are reasonable for (26a), but clearly wrong for (26b): the discourse is incoherent no matter how likely or unlikely such a find is.[9]

These examples motivate an empirical generalization about the behavior of overt disjunctions in the antecedent of a conditional:

(27) **Antecedent Disjunction Entailment**:
If P or Q then R entails both *If P then R* and *If Q then R*.

I don't have a convincing story about what semantic or pragmatic mechanisms generate this entailment. (One possible explanation might refer to the non-cancelable "Need a Reason" implicatures which Lauer (2013) argues are quite generally associated with the use of a disjunction. No doubt a connection with free choice phenomena is also in order.) Fortunately, it's not too important for our purposes to discern exactly how the entailment is generated: what is important is that it is a valid inference, and that it is not enforced by the best available theories of the conditional.

However this entailment is explained, this issue is important for us for two reasons. First, examples involving overt disjunctions in the antecedent of a conditional are frequently used in arguments for the grammatical theory (e.g., by Sauerland (2010); Chierchia et al. (2012); Chierchia (2013)). The fact that (27) survives as a special feature of **overt** disjunctions, despite the general non-monotonicity of conditionals, introduces a major confound into these arguments. In addition, the validity of the principle renders invalid what might appear to be an argument in favor of the grammatical theory. Here is a pair of examples that we saw in §1 (note their similarity to (26)).

(28) a. If it's warm, we'll go outside. But if it's hot we won't go outside.
b. If it's warm or hot, we'll go outside. #But if it's hot we won't go outside.

The first sentence of (28b) clearly entails that we'll go outside if it's hot, and so the continuation is bad.

It might appear that the grammatical theory explains this divergence without reference to an additional principle like (27). This is not the case: once we've adopted a plausible theory of conditionals, silent exhaustification does not make it possible to account for the infelicity of (28b) without also invoking (27). If we have a theory with silent enrichment of *warm*, but without the Antecedent Disjunction Entailment, the situation is exactly as in the cases above involving mutually-exclusive items (such as (25)): whenever it's more likely to be warm-but-not-hot than it is to be hot, the less likely *hot* disjunct in (28b) should be ignored. Thus both discourses should be acceptable and indeed equivalent in such a context.

[9] Chierchia et al. (2012) argue that there is an additional problem involving disjunctions in which one disjunct entails the other, which they call "Hurford's constraint". However, Potts (2013) marshalls a variety of naturally-occurring examples which violate this generalization, arguing that there is no empirical evidence that such a constraint exists. Here I will consider only issues which are directly relevant to the semantic behavior of such disjunctions in the antecedent of a conditional, leaving this broader debate about disjunction to play out on its own.

In any case, the first sentence of (26b) strikes me as unobjectionable, though perhaps somewhat odd if *indeed* is omitted. This is unexpected if Hurford's constraint is operative here.

The grammatical theory would then predict only one difference between the sentences in (28): the (a)-sentence should be ambiguous between a reading which does not give rise to the Epistemic Inference (19) and one which does; but (28b) should give rise to the Epistemic Inference no matter what. Clearly, this is not yet a correct characterization of the data: (28b) is incoherent, no matter what.

In order to render (28b) incoherent, grammatical theorists have two options: return to a downward monotonic theory of conditionals (without von Fintel-style domain-shifting), or adopt a non-monotonic theory together with (27). The first option is out since it saddles us again with invalid inference patterns. We need (27), but once this principle is in place there is no work left for exhaustification to do: (27) alone does the job. Here again, an apparent argument in favor of the grammatical theory turns out to be accounted for by very general features of conditionals, with no need for the additional mechanisms that this theory postulates.

5 Outlook: More confounds, more interactions, but maybe a light at the end

Plausible theories of the semantics of conditionals simply do not support the arguments that proponents of the grammatical theory of scalar implicatures have made. It's not just that there is another way to "avoid the contradiction" that Chierchia (2013) is worried about. There is no contradiction to avoid, since theoretical mechanisms that are needed to deal with very general problems in the semantics of conditionals automatically explain the data that have been adduced in favor of this theory.

Although this means that some of the most influential arguments in favor of the grammatical theory have been misplaced, it does not yet constitute an argument **against** the grammatical theory — except in the weak sense that a theory with fewer mechanisms is *ceteris paribus* preferable. Relative to our current state of understanding of theory and data, the strongest conclusion that we can reasonably draw is that the behavior of scalar items in conditional antecedents is completely uninformative about whether scalar implicatures are generated by Gricean reasoning of by an unpronounced morpheme *exh*.

However, the two theories **do** make distinct empirical predictions, and so it may be that we can squeeze some useful information out of conditional antecedents yet. As we've discussed, cutting-edge theories of the conditional make a prediction about our crucial examples which I called the "Epistemic Inference": *If P then R* and *If P and Q then not-R* jointly entail that *P and not-Q* is more likely than *P and Q*. This prediction also holds with antecedents in a similar semantic subset relationship, such as *warm* ($\approx P$) and *hot* ($\approx P$ *and* Q). However, the predictions of the grammatical theory are less categorical. With the additional mechanisms of exhaustification, conditionals of this form which contain a scalar item are predicted to be ambiguous. On the reading without *exh*, the Epistemic Inference arises; on the reading with *exh*, it does not.

What is less clear is how to operationalize this difference to make it amenable to empirical investigation. Subtle epistemic predictions are probably not best addressed by consulting intuitions, and so convincing evidence on this front will probably require controlled experimentation using methods borrowed from psycholinguistics and the psychological study of reasoning. This effort is, however, likely to be subject to a number of potential confounds.

First, the observation that our grammar generates an ambiguity does not in itself tell us how language users should choose between the interpretations available to them (cf. Potts 2013). Assuming that all sides make the right choice and adopt a good theory of conditionals, Griceans predict one reading, with an epistemic inference. The grammatical theory predicts two, one with an epistemic inference and one without. Whether a verifiable empirical difference emerges depends crucially on our linking assumptions about how the existence of a semantic ambiguity relates to the interpretive choices that listeners make in a particular experimental context.

For example, we might suppose following the Strongest Meaning Hypothesis (SMH; cf. Dalrymple et al. 1998; Winter 2001) that listeners should prefer the meaning which is logically stronger as long as it is consistent with the context. In this case, the logically stronger meaning is the one *without* silent exhaustification, since it generates the Epistemic Inference. Thus, according to SMH, conjunctions of conditionals of the relevant variety should preferentially be interpreted *without* silent exhaustification, *even if the grammatical theory is correct*. If this linking assumption is adopted, the differences between the two theories will only emerge in contexts which specifically contradict the Epistemic Inference. Even then, drawing inferences will be difficult (for Griceans, at least) because Griceans predict no difference, and null results are not interpretable in standard statistical methodology (though see Gallistel 2009; Kruschke 2012). Thus an experiment of this type, interpreted with SMH used as a linking assumption between theory and behavioral data, might be able to provide evidence for the grammatical theory; but it could at most fail to disconfirm the Gricean theory.

Another possibility, suggested by Chemla & Spector (2011), is that in certain contexts participants' responses are generated by counting the number of readings which make a scenario true. If this is right, the theories' predictions about experimental subjects' behavior will differ not categorically but in terms of the *size* of the Epistemic Inference: the effect should be stronger if the Gricean theory is correct, but still present if the grammatical theory is correct. (In fact, the difference in predictions is probably even smaller than this due to additional confounds to be discussed momentarily.) On this interpretation, an experiment which is capable of distinguishing the theories will have to have extremely fine resolution on participants' epistemic states.

A second problem is the following: it's not really true that Griceans always predict only one reading for conditional sentences with scalar items in the antecedent. As discussed in particular by Geurts (2010); Geurts & van Tiel (2013), everyone involved in the debate agrees that it is possible to upper-bound the literal meanings of scalar (and other) expressions to under certain circumstances, including with heavy prosodic focus on a scalar item. The difference between the theories comes down to whether we think of this operation as the basic mechanism for generating scalar implicatures, or as a highly marked operation which requires special discursive or grammatical context, such as heavy prosodic focus.

This point is particular relevant since Chierchia et al. (2012) acknowledge in a footnote that non-monotonic theories of the conditional problematize their argument, but plow on nonetheless; the stated reason for doing so is that "clearly" the crucial examples are acceptable when the Epistemic Inference is known to be false. They imagine a restaurant in which it is known that most customers order both a salad and a dessert, and consider (29) in this context.

(29) If you take salad or dessert, you pay \$20. But if you take both there is a surcharge.

I agree that it is possible to read (29) as a felicitous description of food prices in such a context; but that way is best rendered as (30b), not (30a).

(30) a. If you take SALAD or DESSERT, you pay \$20. But if you take both there is a surcharge.
b. If you take salad OR dessert, you pay \$20. But if you take BOTH there is a surcharge.

Even Geurts (2010) would allow local pragmatic enrichment in the antecedent of (30b); but then we don't expect the Epistemic Inference to arise, since the two antecedents are semantically compatible. The crucial test for the grammatical theory is whether the Epistemic Inference arises with neutral intonation, or even with focus on the disjuncts as in (30a); in these contexts, grammatical theorists predict that the Epistemic Inference should be somewhat weaker than Griceans do (though — as we saw above — this may hold only in certain contexts, and it is not clear how much weaker we should expect the inference to be). Discerning which set of predictions better matches the behavior of interpreters will require careful experimental investigation, with control of the prosodic structure of the sentence.

Third, examples involving overt disjunctions such as those in (30) should be avoided at all costs, because of the additional inferences associated with disjunctions in the antecedents of conditionals (§4). This confound is a major issue, as noted above, since proponents of the grammatical theory frequently invoke such conditionals in support of their thesis. For example, the Antecedent Disjunction Entailment (27) introduces an important confound in the interpretation of (30): the "both" alternative is actually **incompatible** with world knowledge once this entailment is taken into account. That is, a "not both" inference could be generated from the first sentence of (30a) simply by the fact that (a) the sentence entails, given (27), that salad and dessert each cost \$20 separately, and (b) restaurants price food in such a way that you always pay more money if you take more items. Confounds like this are easily avoided by using similarly structured examples not involving disjunction (as long as we take care not to place focus on *an*):

(31) If you take an appetizer, you pay \$20; but if you take more than one appetizer there is a surcharge.

Perhaps this example can be put to work for the grammatical theory; I don't know. Certainly the Epistemic Inference is not obviously inappropriate here. In any case, future work on this topic should take care not to use examples which are confounded in this way.

The final confound that I want to mention is this: examples involving multiple conditionals must be interpreted with great care due to the order effects studied by von Fintel (2001);Moss (2012). Von Fintel (2001) points to contrasts like (32).

(32) a. If you find a butterfly, you should take a note and tell me on Monday. If you find a specimen of *Teinopalpus imperialis*, you should call me immediately.

b. If you find a specimen of *Teinopalpus imperialis*, you should call me immediately. # If you find a butterfly, you should take a note and tell me on Monday.

In failures of Strengthening the Antecedent, there is a strong (perhaps categorical) preference for the more semantically general antecedent to be placed first. The need to account for this restriction is one of the main arguments that von Fintel (2001) adduces in favor of the semantically monotonic, pragmatically non-monotonic semantics that was discussed above. He accounts for the contrast by encoding into a dynamic semantics Lewis' (1979) observation that domains of quantification are easy to expand, but difficult to contract. (However, see Moss 2012 for arguments that the oddity of (32b) is pragmatic in nature and that such examples are acceptable in certain contexts.)

This contrast clearly impacts on the special cases of Strengthening the Antecedent which have (wrongly) been invoked as arguments in favor of the grammatical theory. In fact, von Fintel's observation generates an interesting and previously unnoticed prediction once we add the grammatical theory into the picture. Since exhaustification renders the antecedents mutually exclusive in our running example, the order of the conditionals should not matter if the weaker scalar item is exhaustified. That is, we should be able to tell whether silent exhaustification is available in conditional antecedents by asking whether (33a) and (33b) are equally acceptable.

(33) a. If it's warm, we'll go outside. If it's hot we won't go outside.

b. If it's hot, we won't go outside. If it's warm we'll go outside.

The Gricean theory predicts that (33b) should be bad for the same reason that (32b) is. The grammatical theory predicts, in contrast, that (33b) should have a reading in which *warm* is exhaustified, with the result that (33b) should be fully acceptable. More generally, the order effects noted by von Fintel (2001) should be obviated in all and only contexts in which exhaustification is possible.

For what it's worth, my intuitions go with the Gricean theory: (33b) seems like a strange way to describe your plans for the day, except perhaps with prosodic focus on *warm*. I don't want to make too much of this intuition here — carefully controlled experimental work will be needed to determine whether this is an argument against the grammatical theory, or a welcome prediction. But at least we can end the paper on a positive note: there may, in the end, be an unexpected way to squeeze some useful information about scalar implicatures out of the complex, frustrating, fascinating behavior of conditional antecedents.

References

Adams, Ernest W. 1975. *The logic of conditionals: An application of probability to deductive logic*. Springer.

Adams, Ernest Wilcox. 1998. *A primer of probability logic*. Center for the Study of Language and Information.

Bennett, Jonathan F. 2003. *A Philosophical Guide to Conditionals*. Oxford University Press.

Chemla, Emmanuel & Benjamin Spector. 2011. Experimental evidence for embedded scalar implicatures. *Journal of semantics* 28(3). 359–400.

Chierchia, Gennaro. 2006. Broaden your views: Implicatures of domain widening and the logicality of language. *Linguistic inquiry* 37(4). 535–590.

Chierchia, Gennaro. 2013. *Logic in grammar*. Oxford University Press.

Chierchia, Gennaro, Danny Fox & Benjamin Spector. 2012. The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In Klaus von Heusinger, Claudia Maienborn & Paul Portner (eds.), *Semantics: An international handbook of natural language meaning*, vol. 3, 2297–2232.

Dalrymple, Mary, Makoto Kanazawa, Yookyung Kim, Sam Mchombo & Stanley Peters. 1998. Reciprocal expressions and the concept of reciprocity. *Linguistics and Philosophy* 21(2). 159–210.

Edgington, Dorothy. 1995. On conditionals. *Mind* 104(414). 235. doi: 10.1093/mind/104.414.235.

Egré, Paul & Mikaël Cozic. 2011. If-clauses and probability operators. *Topoi* 1–13. `http://www.springerlink.com/content/t61r3u06227811q8/`.

Égré, Paul & Mikael Cozic. to appear. Conditionals. In Maria Aloni & Paul Dekker (eds.), *Handbook of semantics*, .

von Fintel, Kai & Anthony Gillies. 2010. Must... stay... strong! *Natural Language Semantics* 18(4). 351–383.

Fintel, Kai von. 1999. NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics* 16(2). 97.

Fintel, Kai von. 2001. Counterfactuals in a dynamic context. In Michael Kenstowicz (ed.), *Ken hale: A life in language*, 123–152.

Fox, Danny & Roni Katzir. 2011. On the characterization of alternatives. *Natural Language Semantics* 19(1). 87–107.

Gallistel, CR. 2009. The importance of proving the null. *Psychological Review* 116(2). 439.

Geurts, Bart. 2010. *Quantity implicatures*. Cambridge University Press.

Geurts, Bart & Bob van Tiel. 2013. Embedded scalars. *Semantics and Pragmatics* 6. 9–1.

Goodman, Nelson. 1947. The problem of counterfactual conditionals. *The Journal of Philosophy* 44(5). 113–128.

Grice, H. Paul. 1975. Logic and conversation. In P. Cole & J. Morgan (eds.), *Syntax and semantics 9: Pragmatics*, 41–58. Academic Press.

Grice, H. Paul. 1989. *Studies in the Way of Words*. Harvard University Press.

Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30(6). 669–690.

Kaufmann, Stefan. 2005. Conditional predictions: A probabilistic account. *Linguistics and Philosophy* 28(2). 181–231. doi:10.1007/s10988-005-3731-9.

Kolmogorov, Andrey. 1933. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Julius Springer.

Kratzer, Angelika. 1981. The notional category of modality. In Eikmeyer & Rieser (eds.), *Words, worlds, and contexts*, 38–74. de Gruyter.

Kratzer, Angelika. 1986. Conditionals. In *Chicago linguistics society*, vol. 22 2, 1–15.

Kratzer, Angelika. 1991. Modality. In von Stechow & Wunderlich (eds.), *Semantics: An international handbook of contemporary research*, de Gruyter.

Kruschke, John. 2012. *Doing bayesian data analysis: A tutorial introduction with r and bugs*. Academic Press.

Lassiter, Daniel. 2010. Gradable epistemic modals, probability, and scale structure. In N. Li & D. Lutz (eds.), *Semantics & Linguistic Theory (SALT) 20*, 197–215. CLC Publications.

Lassiter, Daniel. 2011. *Measurement and Modality: The Scalar Basis of Modal Semantics*: New York University dissertation.

Lassiter, Daniel. 2014. *Measurement and Modality: The Scalar Basis of Modal Semantics*. Oxford University Press (to appear).

Lauer, Sven. 2013. *Towards a dynamic pragmatics*: Stanford University dissertation.

Levinson, Stephen C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. MIT Press.

Lewis, David. 1973. *Counterfactuals*. Harvard University Press.

Lewis, David. 1975. Languages and language. *Minnesota studies in the philosophy of science* 7. 3–35.

Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 85(3). 297–315. doi:10.2307/2184045.

Lewis, David. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1). 339–359. doi:10.1007/BF00258436.

McGee, Vann. 1985. A counterexample to modus ponens. *The Journal of Philosophy* 82(9). 462–471.

Moss, Sarah. 2012. On the pragmatics of counterfactuals. *Noûs* 46(3). 561–586.

Potts, Chris. 2013. Conversational implicature: Interacting with grammar. Manuscript, Stanford University. `http://www.stanford.edu/ cgpotts/manuscripts/potts-interacting2013.pdf`.

Sauerland, Uli. 2010. Embedded implicatures and experimental constraints: A reply to geurts & pouscoulous and chemla. *Semantics and Pragmatics* 3. 2–1.

Schlenker, Philippe. 2004. Conditionals as definite descriptions. *Research on language and computation* 2(3). 417–462.

Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in logical theory*, 98–112. Blackwell.

Stalnaker, Robert. 1975. Indicative conditionals. *Philosophia* 5(3). 269–286.

Stalnaker, Robert. 1984. *Inquiry*. MIT Press.

Strawson, Peter. 1986. *If* and $\supset$. In R. E. Grandy & R. Warner (eds.), *Philosophical grounds of rationality*, Clarendon Press.

Winter, Yoad. 2001. Plural predication and the strongest meaning hypothesis. *Journal of Semantics* 18(4). 333–365.

Yalcin, Seth. 2010. Probability Operators. *Philosophy Compass* 5(11). 916–937.

6 Appendix: Proof of the Epistemic Inference for Edgington's (1995) probabilistic theory

Suppose *If P then R* and *If P and Q then not-R* are highly credible. For Edgington (1995), this means that both $p(R|P)$ and $p(\neg R|P \wedge Q)$ are high, where $p(\cdot)$ is a probability measure obeying the usual conditions (Kolmogorov, 1933). We will show that that, as long as both

$p(R|P)$ and $p(\neg R|P \wedge Q)$ are greater than 2/3, the Epistemic Inference follows: $p(P \wedge \neg Q) > p(P \wedge Q)$, i.e., it's more likely that P *and not-*Q is true than it is that P *and* Q is true.

In general, $p(A \wedge B) = p(B|A) \times p(A)$. Thus $p(P \wedge \neg Q) > p(P \wedge Q)$ holds if and only if $p(\neg Q|P) > p(Q|P)$ (canceling $p(P)$ from both sides). Since $p(\neg Q|P) = 1 - p(Q|P)$, what we want to show is equivalent to the condition that $p(Q|P) < .5$.

All probabilities in question are conditioned on P, and so the proof need only refer to the sub-distribution $p(\cdot|P)$. The relevant probabilities are uniquely determined by the probabilities of the members of a four-cell partition over P, representing the possible ways that Q and R could turn out on the assumption that P is true. These are given here along with abbreviations defined for readability.

- $x =_{df} p(Q \wedge R|P)$
- $z =_{df} p(\neg Q \wedge R|P)$
- $y =_{df} p(Q \wedge \neg R|P)$
- $w =_{df} p(\neg Q \wedge \neg R|P)$

Suppose that "high probability" for $p(R|P)$ and $p(\neg R|P \wedge Q)$ is cashed out as "probability greater than 2/3". Since $p(Q|P) = p(Q \wedge R|P) + p(Q \wedge \neg R|P)$, our goal is to show that $x + y < .5$ as long as $p(R|P)$ and $p(\neg R|P \wedge Q)$ are both greater than 2/3.

We can decompose the conditional probability statements of interest into:

$$\begin{aligned} p(R|P) > 2/3 &\quad \text{iff} \quad (x+z)/(x+z+y+w) > 2/3 \\ p(\neg R|P \wedge Q) > 2/3 &\quad \text{iff} \quad y/(x+y) > 2/3 \end{aligned} \tag{1}$$

Using the fact that $x + z + y + w = 1$ and algebraic manipulation, we find

$$\begin{aligned} p(R|P) > 2/3 &\quad \text{iff} \quad x+z > 2/3 \\ p(\neg R|P \wedge Q) > 2/3 &\quad \text{iff} \quad y > 2x \end{aligned} \tag{2}$$

Since $x + z > 2/3$, we also have $y + w \leq 1/3$. Since w is non-negative, $y \leq 1/3$. Given $y > 2x$ and $y \leq 1/3$, $x < 1/6$, and so $x + y < 1/3 + 1/6 = .5$. Since $x + y = p(Q|P)$, we have the desired result: $p(\neg Q|P) > p(Q|P)$, i.e., $p(P \wedge \neg Q) > p(P \wedge Q)$. □