THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Decomposing relevance in conditionals

OPEN ACCESS

WILEY

Decomposing relevance in conditionals

Daniel Lassiter 🆔

Linguistics & English Language,
University of Edinburgh, Edinburgh, UK

Correspondence
Daniel Lassiter, Linguistics & English
Language, University of Edinburgh,
Edinburgh EH8 9AB, UK.
Email: dan.lassiter@ed.ac.uk

Conditionals frequently convey that the antecedent is relevant to the consequent. Recently many authors have argued that this relevance is part of the conventional meaning of conditionals, but this approach fails to account for many examples where a conditional is used to convey *ir*relevance of antecedent to consequent. Both types of conditionals are best explained by a conventional meaning with no relevance requirement, and a separate process of coherence establishment among successive clauses in discourse. This account is supported by the distribution of discourse particles and is able to account for experimental studies used to support the conventionalist position.

KEYWORDS
conditionals, discourse coherence, discourse particles, pragmatics, relevance

1 | INTRODUCTION

The use of a conditional sentence often suggests that the antecedent is relevant to the consequent. Consider (1):

1. If Mary left the party early, Bill was unhappy.

A listener would readily infer that there is some connection between the antecedent "Mary left the party early" and the consequent "Bill was unhappy." They do not, for example, merely happen both to be true. This connection does not determine a direction of causation: (1) is compatible with a situation where Mary's early departure caused Bill to feel unhappy, or one where the

only explanation for Mary's departure is that she loathes being around an unhappy Bill. With the right background, it may also be compatible with a "common cause" scenario in which some third variable generates a correlation between the two events mentioned in (1). Whatever the causal story, this example clearly conveys that the question of whether the antecedent is true is *relevant* to the truth-value of the consequent. Following Skovgaard-Olsen et al. (2016), let us call this the *relevance effect*.

While the relevance effect is widely known, most theoretical accounts of conditionals have not built a relevance requirement into the literal meaning of conditionals, and it is widely assumed to be a conversational implicature. However, this claim is rarely justified in any detail. Recently a number of authors have given strong theoretical and empirical arguments showing that the relevance effect does not behave like a conversational implicature (Douven, 2008, 2016; Krzyżanowska, 2019; Krzyżanowska & Douven, 2018; Skovgaard-Olsen, 2016; Skovgaard-Olsen et al., 2016). These authors have generally concluded, by process of elimination, that the relevance requirement is part of the conventionally encoded meaning of conditionals, either as an entailment or a conventional implicature. This position has been dubbed "inferentialism" (Douven et al., 2018). Inferentialism marks an important departure from all major philosophical, linguistic, and psychological theories of conditionals, including those based on the material conditional, possible-worlds semantics, domain restriction, mental models, suppositional reasoning, or probability logic.

This article proposes an alternative non-conventionalist approach to the relevance effect, rooted not in conversational implicature but in *discourse coherence*. Drawing on a rich body of work on discourse coherence (Asher, 1993; Asher & Lascarides, 2003; Hobbs, 1979, 1990; Kehler, 2002; Knott & Dale, 1994, 1996), I suggest that the need to establish relevance between antecedent and consequent in examples like (1) is due to an obligatory process of inferring coherence relations among successive pieces of discourse. Coherence effects differ from conversational implicatures in numerous ways; arguments that motivate inferentialism by pointing to the failure of an implicature-based account thus do not count against a coherence theory. Additional support for the coherence account comes from the fact that the relevance requirement is systematically missing in certain conditionals, a fact which requires inferentialist theories to treat *if* as ambiguous between an interpretation that encodes relevance and one that does not. In contrast, the coherence-based theory is able to predict precisely where the relevance effect should occur based on parallels with the context-sensitivity of other multiclause discourses and the effects of discourse particles such as *then* and *still* in conditional and non-conditional contexts.

The issue of whether relevance is semantically encoded in conditionals is crucial for understanding what conditionals mean and how they are used—topics of broad interest in psychology, linguistics, philosophy, computer science, and other fields. The viability of non-inferentialist theories of conditionals—including those based on mental models, supposition/the defective truth-table, probability logic, the strict conditional, the material conditional, and the Stalnaker conditional—may also depend on it. In addition, the status of the relevance effect matters for several very general issues in the psychological and philosophical study of conditionals.

For instance, a conventionally encoded relevance requirement could—depending on how it is spelled out in formal detail—problematize *strong centering*, a property of many popular theories of conditional semantics:

> **Strong centering**: If A and C are both true, then "If A, C" is true.

The status of strong centering is an important issue in philosophical logic, and it also enjoys some experimental support (Cruz et al., 2015; Politzer & Baratgin, 2016). However, it is

problematic if the relevance effect is part of the semantics of conditionals, since the mere truth of A and C does not suffice to guarantee relevance between them (Cruz et al., 2016; Krzyżanowska & Douven, 2018). The status of the relevance effect is also important because it bears directly on the question of whether the probability of a conditional is systematically equal to the corresponding conditional probability (Adams, 1975; Edgington, 1995; Lewis, 1976; Stalnaker, 1970).

Stalnaker's thesis: P(If A, C) = P(C|A).

This equation has been supported in a large number of empirical studies (Douven & Verbrugge, 2010, 2013; Fugard et al., 2010; Hadjichristidis et al., 2001; Over & Evans, 2003). However, a semantic relevance requirement would generate systematic counter-examples to Stalnaker's thesis, since irrelevance of A to C could lead to low probability of "If A, C" even if the conditional probability P(C|A) is high (Cruz et al., 2016; van Rooij & Schulz, 2019). However, an account of the relevance effect in terms of discourse pragmatics is able to explain this result without abandoning Stalnaker's thesis, using similar logic to the account of strong centering (Section 8). The status of the relevance effect is thus one of considerable significance for these and various other issues in the study of conditional semantics, pragmatics, and reasoning.

2 | RELEVANCE AS A CONVERSATIONAL IMPLICATURE?

The basic relevance effect was illustrated above: When we read a conditional such as (1) ("If Mary left the party early, Bill was unhappy") we automatically assume that the antecedent and consequent clauses are informationally related. Many authors writing on conditionals have noticed this fact and suggested that it is due to conversational implicature—though few have attempted to spell out the reasoning in detail. One exception is Quine (1965, Section 7), who argues that "if" denotes the material conditional. However, Quine continues, someone who believes (1) simply because she believes its antecedent to be false "could save breath and at the same time convey more information" by asserting the negation of the antecedent. Similarly, someone who believes (1) because she believes the consequent to be true would be better served by asserting its briefer, more informative consequent. This leaves only one obvious motivation for asserting the conditional: a belief that "there is some causal connection or general law" whereby the antecedent's truth would tend to lead to the consequent being true as well. Quine points out that, if such an account is able to explain the relevance inference without encoding it conventionally, we should resist the temptation to make it part of the meaning of conditionals, as modern inferentialists do.

In an extended defense of the material conditional analysis against an early form of inferentialism, Grice (1989a, written in 1967) echoes Quine's pragmatic explanation of why the assertion of a conditional leads to the inference that the speaker does not merely believe the antecedent to be false, or the consequent true. By reasoning about the conversational maxims of *quantity* and *quality* (Grice, 1989b), we can infer from the assertion of a conditional that there is "non-truth-functional evidence" for the truth of the conditional (Grice, 1989a, pp. 61–62).

The main problem for this account is that the relevance effect in conditionals does not behave like a conversational implicature, according to Grice's own criteria. In particular, as Krzyżanowska (2019) points out, the relevance effect does not appear to be defeasible as

conversational implicatures typically are. Consider the case of canceling a quantity implicature. While the (a) example would typically be enriched to (b), there is no incoherence in the (c) sentence, which explicitly denies the "not all" enrichment.

2. a. Many Americans own flags.
   b. Many, but not all, Americans own flags.
   c. Many Americans own flags—indeed, all of them do.

In contrast, an attempt to cancel the relevance effect explicitly generally leads to a sense of bizarreness, as in (3).

3. If Mary left the party early, Bill was unhappy—though these things have nothing to do with each other.

Work by Skovgaard-Olsen et al. (2019, Experiment 1) corroborates the intuition that relevance in conditionals differ from scalar implicatures. Their participants rated the denial of relevance in a conditional as much more contradictory than the denial of a default scalar implicature, and indeed as contradictory as the denial of an uncontroversial entailment.

A related problem for an implicature-based story is the phenomenon of "missing-link conditionals." The requirement of a connection between antecedent and consequent appears to be obligatory in indicative conditionals (Douven, 2017). In a context where such a connection is implausible or difficult to recover, the example is felt to be conversationally inappropriate.

4. If Mary left the party early, Bill is Swedish.

Example (4) is quite difficult to make sense of, apparently because it is difficult for a listener or reader to supply a relevant connection between the timing of Mary's departure and Bill's nationality. Of course, it is possible to make sense of (4) with further contextual background—for example, if we know that Mary has an irrational dislike of Swedes (and only Swedes). In this case, knowledge of her departure would allow us to infer Bill's nationality, rendering (4) sensible.

The need to supply such unusual background in order to satisfy the relevance requirement in (4) sets it apart from classic conversational implicatures such as scalar implicatures. Specifically, in the original Gricean account as well as modern instantiations (Chierchia et al., 2012; Frank & Goodman, 2012), the fact that an implicature is very implausible would normally lead to the suspension of the implicature, rather than a sense of bizarreness. This is clear if we consider variants of example (2a) in which the "many" $\rightsquigarrow$ "not all" implicature would be implausible or nonsensical: no infelicity results, and there is simply no implicature.

5. a. There are many black holes in the universe.   (?? "There are not all black holes")
   b. The hotel has many amenities.   (?? "It does not have all amenities")

In contrast, the fact that the antecedent and consequent of (4) are irrelevant to each other in a typical context does not lead to the suspension of the relevance requirement, but rather to a generalized sense of inappropriateness. The need to supply special context to render (4) felicitous suggests that establishing relevance in conditionals is obligatory, rather than context-sensitive and defeasible. As Douven (2017), Krzyżanowska (2019), and others have pointed out, this suggests that the relevance effect is not a conversational implicature.

A number of other arguments have been put forward in recent literature to problematize an implicature-based account. For example, Krzyżanowska and Douven (2018) provide experimental evidence that participants respond differently to failures of relevance in conditionals than to false implicatures. Krzyzanowska et al. (2017) show that participants reject "If *p*, *q*" conditionals when *p* and *q* are true but irrelevant to each other, while accepting "*p* and *q*"; but this pattern is reversed when *p* and *q* are inferentially relevant. The latter point is particularly problematic if quantity-based reasoning is involved in the generation of the relevance effect, since on such an account "*p* and *q*" should be preferred as the stronger option. Additional arguments against the implicature-based account are given by Douven (2008, 2015); Skovgaard-Olsen (2016); Skovgaard-Olsen et al. (2016); and Skovgaard-Olsen, Collins, et al. (2019).

3 | RELEVANCE AS CONVENTIONALLY ENCODED?

There seems to be good reason to reject an initially attractive Gricean account of the relevance effect, then. What is the alternative? Inferentialist theories are motivated in large part by the failure of the implicature-based account, reasoning that the only alternative is for relevance to be a conventionally encoded part of the meaning of conditionals (Douven, 2008; Krzyżanowska & Douven, 2018; Skovgaard-Olsen, 2016, among others). As Douven et al. summarize the position:

> [I]nferentialism is the only semantics that makes the existence of an inferential connection between antecedent and consequent a requirement for the truth of a conditional. Inferentialism, in other words, builds the requirement of a connection into the meaning of the word "if", thereby straightforwardly accounting for the felt oddness of [missing-link] conditionals. (Douven et al., 2018, p. 51).

This formulation is noncommittal about the status of the relevance requirement as an entailment, presupposition, conventional implicature, or some other type of conventional meaning. However, Skovgaard-Olsen, Collins, et al. (2019) find stark differences between the relevance effect on the one hand and the behavior of semantic presuppositions and conversational implicatures on the other. They suggest that the relevance effect most resembles the type of conventional implicature that Grice (1989b) identifies in the particle "therefore." Skovgaard-Olsen et al. (2016) also gives a theoretical analysis connecting the inferentialist approach to *if* with "therefore," "but" and similar items that appear to carry conventionally encoded meaning that is not at issue but also not presuppositional.[1]

While inferentialism does account straightforwardly for the robustness of the relevance effect when it occurs, it may do this job too well. For there are numerous examples of conditionals that are *not* associated with a relevance effect, like (6).

6. If Mary leaves now, Bill will still be unhappy.

The message of (6) is precisely the opposite of what inferentialist theories would lead us to expect: It conveys that the consequent will be true regardless of the status of the antecedent, that is, that antecedent and consequent are *ir*relevant to each other.

---

[1]In fact a full theory of the conversational effects of "therefore" and "but" would naturally be framed in terms of their effects on discourse coherence: "therefore" marks *result* and "but" marks *violated expectation* (see Section 4). In a sense, then, Skovgaard-Olsen's argument foreshadows the positive account proposed here.

The addition of *still* in (6) is helpful in generating the irrelevance reading of conditionals, and I will explain why in Section 4 below. But irrelevance readings can arise without such additional material, as in (7): neither conditional can be associated with a relevance effect, on pain of contradiction.

7. a. Q: If Mary leaves now, will Bill be unhappy?
   b. A: If Mary leaves now, Bill will be unhappy. If she doesn't leave, he'll be unhappy. It makes no difference.

The inferentialist account of "if" would predict that both conditionals in (7b) should conflict with the continuation "It makes no difference." However, A's response is perfectly acceptable in this context, and conveys precisely that Mary's departure is irrelevant to Bill's unhappiness.[2]

How can we account for the fact that the use of a conditional sometimes conveys that the antecedent and consequent are relevant to one another, and sometimes that they are irrelevant to one another? The usual response seems to be that the relevance effect holds for "normal" uses of conditionals, and that examples like (6) and (7) can be treated as somehow non-standard. However, the notion of a "normal" use does not have a theoretical interpretation in natural language semantics. If we dismiss a particular use of an item as "abnormal," this amounts to postulating an ambiguity. Douven embraces this consequence explicitly:

> As has been frequently observed in the literature, the word "if" is put to a great many uses ... Some might still hope for a unified theory of conditionals that correctly predicts all possible judgments about any conditional any competent user of the language might ever make. Note, though, that it is by no means a priori that such a theory can be had. ... [I]t seems more realistic to suppose that the right approach to conditionals is of a more modest and piecemeal variety. In this spirit, I offer the evidential support theory of conditionals as a theory of ordinary or normal uses of conditionals. "Normality" is here to be taken in the entirely unproblematic statistical sense of the word. (Douven, 2008, p. 31).

Similar diagnoses of conditionals that fail to display the relevance effect are given by, for example, Douven (2017) and Skovgaard-Olsen et al. (2016). These authors vary in whether they treat the ambiguity as involving a lexical ambiguity in the meaning of "if" or an ambiguity in logical form, for instance, the presence of a silent *even*.

A natural worry is that this type of analysis deprives the inferentialist approach to conditionals of explanatory value: the theory seems now to be that sentences with "if" encode a

---

[2]A reviewer suggests treating (7b) as concessive conditionals with no relevance requirement. As I explain below, this is a theoretically unattractive move since it requires postulating unnecessary ambiguity. In addition, if one of the conditionals in (7b) is concessive and the other is "normal," then inferentialism continues to predict that the dialog should be incoherent due to conflict between the relevance requirement of the normal "if" and the continuation "It makes no difference." If both conditionals are concessive, the sequence is unintelligible:

1. *Even* if Mary leaves now, Bill will be unhappy. *Even* if she does not leave, he'll be unhappy. It makes no difference.

Even with considerable freedom in postulating otherwise unmotivated ambiguities, there is no way for inferentialism to account for the felicity of (7b).

relevance requirement, except when they do not. This objection would be softened if inferentialists were to provide guidance about which reading should be selected in which contexts. However, no one has attempted to do so, beyond the very general claim that the relevance-enforcing reading is somehow "normal" or "default." It would clearly be more satisfactory to have a theory of conditionals that provides a unified interpretation for "if" in both relevance-supporting and relevance-denying readings, and an account of which contexts support which readings. In other words, the theoretical situation would be greatly improved if we had a theory that explains why statistics of use seem to favor the relevance requirement—and treats less frequent uses not as anomalies to be ignored, but as important clues to the nature of the relevance effect. The coherence-based account presented below attempts to do precisely this.

4 | RELEVANCE AS A COHERENCE EFFECT

In order to account for both relevance and irrelevance readings of conditionals, inferentialism requires us to postulate a semantic ambiguity in conditionals. However, an independently motivated theory of discourse interpretation generates the observed interpretations without any need to postulate ambiguity.

I propose that the relevance effect can be decomposed into two components: a basic semantics for conditionals that is blind to relevance, and an obligatory pragmatic process of establishing *coherence relations* among adjacent clauses in discourse (Asher, 1993; Asher & Lascarides, 2003; Hobbs, 1979, 1990; Kehler, 2002; Knott & Dale, 1994, 1996). This allows for the frequent need to establish relevance between antecedent and consequent, and also for the possibility of using a conditional to indicate irrelevance. As we will see, neither use is a special feature of conditionals. Both are instances of a general interpretive process whereby we infer coherence relations among successive pieces of discourse, and are attested in a variety of other multiclause texts.

4.1 | Ambiguity or discourse pragmatics?

At a high level, the coherence-based account of the relevance effect that I will propose echoes Grice's (1989b) defense of the Boolean interpretation of connectives such as *and*. On face, *and* is used in a bewildering variety of ways, a few of which are illustrated in (8).

8. a. Venus is smaller than Earth, and Mercury is smaller than Venus.   ("and similarly")
   b. The clock reached 6 AM, and the rooster crowed.   ("and then")
   c. Bill laughed at John, and John slugged him.   ("and as a result")
   d. Bill got caught embezzling money, and he's a judge!   ("and yet")

We might suppose that the word *and* is several ways ambiguous, with a similarity reading, a temporal reading, a causal reading, and a contrastive reading, perhaps among others.

As with the inferentialist position on conditionals, this idea could be fleshed out in two ways: either *and* is multiply lexically ambiguous, or there are silent operators—a silent *then* in (8b), and so on. Grice (1989c) argues that we should not posit such ambiguities based merely on the existence of multiple intuitive interpretations, unless we can rule out plausible alternative accounts that could allow us to retain a unified meaning.

> **Grice's modified Ockham's razor.** Senses are not to be multiplied beyond necessity.

Grice's theory of conversational implicature provides a device for enriching the meanings of sentences to include additional information derived on the basis of general reasoning processes. He suggests that this sort of reasoning could derive at least some of the intuitive readings of *and*. If a theory along these lines is successful, the motivation for postulating multi-way ambiguity in the meaning of *and* disappears: This was reasonable only when we were unable to generate the needed enrichments via general and independently motivated mechanisms. Crucially, the Gricean position requires that *and* be given a skeletal interpretation with no information about time, causation, contrast, and so forth, which is then enriched by general pragmatic mechanisms.

While Grice's methodological point is sound, his suggestion to use conversational implicature to explain examples like (8) has not stood the test of time. Subsequent work in discourse pragmatics has shown that the various "readings" of *and* in (8) can be attributed to varying *coherence relations* among clauses in discourse. Crucially, these relations are not directly tied to the item "and," but are widely attested across other connectives and also in juxtapositions of sentences and in multispeaker discourses. For example, if we convert each example in (8) to a juxtaposition of two sentences, the same intuitive readings emerge.

9. a. Venus is smaller than Earth. Mercury is smaller than Venus.   ("Similarly, ...")
   b. The clock reached 6 AM. The rooster crowed.   ("Then, ...")
   c. Bill laughed at John. John slugged him.   ("As a result, ...")
   d. Bill got caught embezzling money. He's a judge!   ("Astonishingly, ...")

The *exact same interpretations* occur in the matched texts with and without *and*. Given this, it would be theoretically extravagant to postulate an ambiguity in the meaning of *and* to account for (8). What we have instead are a variety of ways of relating the content of successive clauses to each other, which occur in various places in discourse—including conditionals.

In what follows I rely largely Kehler's (2002) theory of coherence, which is adapted from that of Hobbs (1979, 1990). This is primarily for simplicity of presentation: the main points are compatible with other major theories of coherence in the literature.

4.2 | Result and explanation

Consider the text in (10).

10. Mary left the party early. Bill was unhappy.

This text would not normally be interpreted as a string of unconnected utterances, but rather as involving some kind of informational relationship between the sentences. One natural interpretation invokes a cause–effect relationship: Mary left early, and *as a result* Bill was unhappy. This is an instance of the *result* relation of Kehler (2002):

> **Result**: Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally
> P $\leadsto$ Q.   (Kehler, 2002, p. 20)

The "⤳" in this definition is glossed by Kehler as "could plausibly follow from"; it is meant to be weaker than the conditional.

The *result* relation is not explicitly marked in (10), and has to be inferred. It can be marked explicitly by various particles and adverbials:

11. Mary left the party early. So/As a result, Bill was unhappy.

Result relations do not only occur in multisentence utterances, but also in multispeaker discourses and within single sentences. For example:

12. a. **Speaker A**: Mary left the party early. **Speaker B**: (So) Bill was unhappy (then).
    b. Mary left the party early, (and) so Bill was unhappy.

Most crucially for us, the *result* relation also occurs naturally in conditionals and other subordinating contexts that involve suppositional reasoning.

13. a. If Mary left the party early, Bill was unhappy (as a result).
    b. Assuming that Mary left the party early, Bill was unhappy (as a result).
    c. Supposing that Mary left the party early, Bill was unhappy (because of her departure).

The other cause-effect relation discussed by Kehler (2002) is *explanation*, which is simply *result* with the order of clauses reversed.

Explanation: Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally Q ⤳ P.   (Kehler, 2002, p. 21)

An explanation relation is also possible in the bare juxtaposition of sentences in (10), which can be read as suggesting that Mary left early *because* Bill was unhappy. This reading comes out clearly in examples where a result relation is implausible, such as (14).

14. Mary left the party early. The host was drunk and abusive.

Explanation relations also occur in multi-speaker discourses and various inter-sentential contexts, for instance:

15. a. Speaker A: Mary left the party early.
       Speaker B: The host was drunk and abusive.
    b. Mary left the party early, because the host was drunk and abusive.

Like result, explanation relations are also possible in conditionals: Both examples in (16) indicate that we can infer the existence of the cause described in the consequent from the effect that is considered suppositionally in the antecedent.

16. a. Supposing/Assuming that Mary left the party early, the host was drunk and abusive.
    b. If Mary left the party early, the host was drunk and abusive.

What sort of meanings do result, explanation, and other coherence relations provide? Since they can occur with no linguistic trigger in simple juxtapositions, coherence relations are the product of discourse pragmatics, rather than being conventionally associated with material in the linguistic forms of sentences. Indeed, Kehler (2002) provides a careful explanation of how the particular relations that he postulated arise from domain-general reasoning about possible connections among ideas.

Coherence relations do not appear to modify truth-conditional content. For instance, a speaker who utters (14), knowing that Mary left early because she needed to catch a train, has surely said something misleading. He has not, however, said anything false: His utterance was the juxtaposition of two true assertions, after all. In this sense, coherence relations may appear to enrich the conveyed meaning of a text beyond its semantic content in much the same way that conversational implicatures do. However, not all coherence relations can be analyzed in this way. For instance, Kehler's *exemplification* and *parallel* relations are purely discourse-internal, "signposting" how the information in a clause is meant to be situated with respect to the broader structure of the discourse and its goals. This sort of meaning is related to discourse topic and other "information packaging" concepts that play a critical role in managing conversational structure and ensuring successful communication (Chafe, 1987).

This type of meaning can be understood from within a variety of competing theoretical frameworks. In Kehler's theory, coherence relations are essentially presuppositional, constraining the worldly and conversation–internal informational background of the discourse. The presuppositions that they carry are, however, quite readily accommodated in the sense of Lewis (1979). In other accounts relations among units of text might be given a more explicitly interactive meaning (Clark, 1996; Schiffrin, 1987). It is not necessary to choose among these high-level perspectives here. However, it is important that coherence relations may well involve *sui generis* conversational or interactive meaning components that defy efforts to force them into the propositionally-oriented meaning categories familiar from Gricean pragmatics.

4.3 | Inferential particles

So far we have seen two very general coherence-establishing strategies—*result* and *explanation*—both of which occur in conditionals among other multiclause discourses. These are two straightforward ways for two pieces of information to be in a relevance relation: A and C are relevant to each other whenever A typically allows one to infer C, or C typically allows one to infer A. These relations can be inferred, or can be provided explicitly by a marker that picks out one of these relations: *as a result* singles out a result relation, while *because* singles out explanation. Some particles are more general, indicating an inferential relationship between clauses without constraining what kind of inference. Two notable examples are *then* and *so*, both of which are compatible with inferential connections even when there is no causal relationship in either direction, but both are correlated via a third variable.

The cook would have used a knife as the murder weapon, and would have hidden it in the pantry; the butler would have used a cane, and would have hidden it in the garden shed. Neither the weapon's location nor its identity is a cause of the other, but in this context either can be used to infer the other.

17. a. A: The weapon was a knife.
    B: So/Then it was found in the pantry.
    b. The weapon was a knife, and so it was found in the pantry.
    c. The weapon was found in the pantry. So, it was a knife.

In these examples, *so* and *then* function to indicate that the information in the second clause can be inferred from the first, without commenting on what the precise shape of the causal relationship is.[3]

These points may shed light on the vexed question of conditional *then*. While *if ... then* has often been treated as a kind of discontinuous operator that is synonymous with bare *if*, many have noted that this analysis is both linguistically implausible and non-explanatory (e.g., Grice, 1989a, p. 62). On the influential analysis of Iatridou (1993), the inference from *p* to *q* is part of the meaning of *if*, but *then* adds a presupposition that *q* cannot be inferred from ¬*p*. The effect of *then* is that we cannot infer the consequent no matter what, ensuring that the antecedent is relevant to our ability to infer the consequent.

Combining elements of Grice's and Iatridou's analyses within a discourse-theoretic framework, Biezma (2014, p.374) argues convincingly that "*then* is a discourse marker that ... establishes an anaphoric relation between the information gained from the previous discourse move and information gained from the subsequent move." As one instance of this discourse function, *then* in conditionals functions to enforce a "causal explanatory" relation between antecedent and consequent. It plays a similar role in other types of sentences and texts, as we have seen (see Biezma's article for much more linguistic detail and analysis).

Biezma's account is highly compatible with the coherence-based theory pursued here: a "causal explanatory relation" will generally involve result or explanation.

18. a. If Mary left the party early, then Bill was unhappy (as a result).
    b. If Mary left the party early, then the host was drunk and abusive (we can infer).

This line of analysis predicts that conditionals with *then* should be incompatible with contexts where there is no inferential relationship between antecedent and consequent. Consider the following pair due to Davis (1983):

19. a. If you open the refrigerator, it will not explode.
    b. If you open the refrigerator, then it will not explode.

Given our normal assumptions about refrigerator behavior, (19a) is naturally read as an irrelevance conditional: the refrigerator will not explode tout court, including in the case where you open it. However, example (19b) with *then* seems to exclude this interpretation. As Davis notes, it conveys the surprising claim that opening the refrigerator will somehow prevent it from exploding, with the disturbing implication that an unopened refrigerator will explode. This pattern is explained if we suppose that the conditional is semantically compatible with both relevance and irrelevance readings, but the use of *then* adds a discourse-based relevance requirement due to its association with result and explanation.

The discussion of particles and connectives so far has focused on what they are used to *do*, in terms of their correlations with discourse relations. But what do they *mean* that allows them to have these functions? There are many perspectives on this issue (see Schiffrin, 2001 for an overview). However, in terms of the connections with coherence relations that are explored here, a few styles of analysis suggest themselves for particular items.

---

[3]One oddity is that the combination *and then* does not have an inferential reading: It appears to be compatible only with the temporal meaning of *then*.

The first is that some items may conventionally pick out coherence relations. This is perhaps most plausible for items like "likewise," "nevertheless," and "however," which do not make a truth-conditional contribution but indicate something about how the upcoming linguistic material is meant to relate to a previous discourse move. Some particles cannot be treated this way: for instance, "then" is compatible with both result and explanation. This suggests that certain items may contribute presuppositions or conventional implicatures whose effect on coherence relations is indirect: A coherence relation that conflicts with a sentence's presupposition or conventional implicature is ruled out, leaving only those that are consistent with it. It is not clear to me whether these two positions are in competition. Perhaps discourse-anaphoric meanings can be seen as a special case of conventional implicatures, for instance.

While "then" and "nevertheless" have discourse-related, non-truth-conditional meanings, some coherence-related markers have truth-conditional effects. Consider (20), which is false—not just misleading—if Mary laughed for some unrelated reason.

20. Mary laughed because Bill's socks were unmatched.

Here *because* limits us to the explanation relation because its truth-conditional contribution is semantically incompatible with other available coherence relations. There appear, then, to be several distinct ways that particles and connectives can constrain coherence relations.

4.4 | Irrelevance conditionals and violated expectation

The discussion of conditionals with and without *then* in the last section brings irrelevance readings back to the fore. Recall that our main criticism of inferentialism above was that the theory is unable to give an explanatory account of the variable use of conditionals to indicate either relevance or irrelevance between clauses. The decomposition of relevance effects into conditional semantics and coherence relations, on the other hand, is perfectly suited to do this: irrelevance between clauses is associated with another type of coherence relations, with their own linguistic markers.

The relevant coherence relations here are *Violated expectation* and its reverse-order counterpart *Denial of preventer*.

Violated expectation: Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally P $\rightsquigarrow$ ¬Q. (Kehler, 2002, p. 21)

Denial of preventer: Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally Q $\rightsquigarrow$ ¬P. (Kehler, 2002, p. 21)

Violated expectation is available in a variety of contexts, though it is often more natural with support by particles such as *but*, *yet*, *nevertheless*, *anyway*, *however*, or *still*. For instance, (21) can be read in this way, but the intended reading is more prominent in (22).

21. Mary left the party in a huff. The host was gracious to her as she departed.

22. a. Mary left the party in a huff. Still/Yet/However, the host was gracious to her as she departed.
    b. Mary left the party in a huff. The host was still gracious to her as she departed.

In multi-clause sentences with *and*, a similar pattern emerges: The violated expectation reading is possible with a bare conjunction, but much more salient with support from an appropriate particle as in (24).

23. Mary left the party in a huff, and the host was gracious to her as she departed.
24. a. Mary left the party in a huff, and <u>still/yet</u> the host was gracious to her as she departed.
    b. Mary left the party in a huff, and the host was <u>still</u> gracious to her as she departed.

Violated expectation can also be marked by substituting *but* for *and*. While *and* places few constraints on appropriate coherence relations, the choice of *but* restricts the coherence relation to a handful of relations—one of which is violated expectation.

25. Mary left the party in a huff, but the host was gracious to her as she departed.

Precisely the same pattern occurs in conditionals. It is difficult to construe the conditional in (26) as indicating a violated expectation relation. The most obvious reading involves *result*, which is surprising given that the content does not readily support this reading in light of our knowledge about typical human behaviors.

26. If Mary left the party in a huff, the host was gracious to her as she departed.

However, with support from appropriate linguistic marking a violated expectation reading emerges.

27. a. If Mary left the party in a huff, the host was <u>still</u> gracious to her as she departed.
    b. <u>Even if</u> Mary left the party in a huff, the host was gracious to her as she departed.

Note here that *even if* conditionals are naturally used to mark violated expectation relations. This makes sense in light of theories of *even if* such as those of Lycan (2001) and Bennett (2003), which compositionally derive the fact that "Even if A, C" often conveys that "If A, C" and "If not-A, C" are both true.[4]

Denial of preventer relations exhibit a similar pattern, tending to be less prominent unless marked explicitly. With the noncommittal connective *and* the reading does not emerge readily, but it does with *even though* and *despite*, which (as Kehler, 2002, p. 21 notes) can be used to mark denial of preventer.

28. The host was gracious to Mary as she departed, and she left the party in a huff.

29. a. The host was gracious to Mary as she departed, <u>even though</u> she left the party in a huff.
    b. The host was gracious to Mary as she departed, <u>despite</u> her leaving the party in a huff.

---

[4]In these accounts, *even* has the same meaning that it does in *Even SUE left*. Roughly, the latter conveys that *Sue left* is not only true, but quite unlikely/surprising relative to true focus alternatives such as *Bill left, Sally left*, and so forth. Similarly, a conditional of the form "Even if A, C" typically conveys that "If A, C" is not only true, but quite unlikely/surprising relative to some true alternatives. If the alternative happens to be "If not-A, C"—as it will in many contexts—the total effect is that C holds unconditionally. The relative unlikelihood/surprisingness implication of "Even if A, C" is of course highly compatible with the Violated Expectation implication that one would normally infer from A to ¬C.

Once again, this pattern also holds in conditionals: The bare conditional in (30) naturally favors an explanation relation despite the oddness of this interpretation given world knowledge. However, the corresponding *even if* conditional unambiguously favors a denial of preventer reading.

30. The host was gracious to Mary as she departed if she left the party in a huff.

31. The host was gracious to Mary as she departed <u>even if</u> she left the party in a huff.

I will suggest below that the apparent default status of result/explanation relations in conditionals is attributable to the dispreference for violated expectation and denial of preventer readings when there is no supporting linguistic marking, across all of these constructions. This, in turn, explains the "normal" status of conditionals that support the relevance effect.

Conditionals whose coherence relation is violated expectation or denial of preventer are generally interpreted as irrelevance conditionals. For instance, our initial example of an irrelevance conditional in (6) involved a violated expectation relation marked by *still*.

32. If Mary leaves now, Bill will still be unhappy.

Why do we infer from (32) that Bill will also be unhappy if Mary does not leave—and so that Bill will be unhappy no matter what, rendering antecedent and consequent irrelevant to each other? According to Kehler's definition, a violated expectation relation here indicates that there is a default expectation that one can infer from the antecedent to the negation of the consequent. In this case, that means that we would normally expect Mary's departure not to be associated with Bill's unhappiness in a situation like the one described. However, this condition is still compatible with a set of default expectations in which Bill will *also* be unhappy if Mary does *not* leave—that is, where our prior expectation, before hearing (32), is that its consequent will be true no matter what.

In other words, Kehler's definition of violated expectation predicts that violated expectation relations can occur when background expectations include both $P \rightsquigarrow \neg Q$ and $\neg P \rightsquigarrow \neg Q$. However, *still*, *yet*, and perhaps violated expectation relations in general are not appropriate in such a context. For example, consider (33). In general, we expect both the choice to drink wine or not to be irrelevant to dying: both **wine** $\rightsquigarrow \neg$**death** and $\neg$**wine** $\rightsquigarrow \neg$**death**. On Kehler's definition, violated expectation should be possible in such a context, and the conditional in (33) should convey that the wine is poisoned.

33. If you drink this wine, you will still die.

But (33) does not have this interpretation: it seems appropriate only in an unusual context where it is expected that you will die if you do *not* drink the wine, and the wine is meant to *save* you from death. In such a special context, (33) would indeed convey that the wine will not have the desired effect of saving you from death.

The fact that such an odd context is needed to render (33) acceptable suggests that Kehler's conditions for Violated Expectation need to be strengthened. This relation is not appropriate when there is a background expectation that the second clause (here, the conditional consequent) will be false no matter what.

We can account for these observations in either of two ways: by strengthening the definition of violated expectation, or modifying the meaning of the relevant particles. If we choose the former route, the necessary modification seems to be:

Violated expectation (modified): Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally $P \rightsquigarrow \neg Q$ and $\neg P \rightsquigarrow Q$.

Denial of preventer (modified): Infer P from an assertion of S_1 and Q from an assertion of S_2, where normally $Q \rightsquigarrow \neg P$ and $\neg Q \rightsquigarrow P$.

These changes are also motivated by examples not involving conditionals. For instance, (34) is appropriate only under the unusual background assumption that the wine was somehow expected to *save* Bill from death, and not under the typical expectation that drinking wine and dying are irrelevant.

34. Bill drank a glass of wine. He still died.

The odd background required to make sense of (34) follows from the modified definition as follows: *still* marks violated expectation, which is only appropriate when both (a) drinking wine is associated with not dying, and (b) not drinking wine is associated with dying. The latter generates the prima facie oddness of (34), since typical conversational backgrounds do not have this feature.

With this small modification to Kehler's theory, we can explain why conditionals whose coherence relation is violated expectation or denial of preventer are interpreted as irrelevance conditionals. Since *still* marks violated expectation, (35) is only appropriate when both (36a) and (36b) are in the context as default assumptions.

35. If Mary leaves now, Bill will still be unhappy.

36. a. If Mary leaves now, Bill will not be unhappy.
    b. If Mary does not leave now, Bill will be unhappy.

Since the assertion (35) contradicts the default (36a), the latter assumption is discarded. However, (36b) is consistent with the assertion and remains. The resulting context contains the information in assertion (35) and assumption (36b). The net effect is a context that supports the following:

37. If Mary leaves now, Bill will be unhappy; and if Mary does not leave now, Bill will be unhappy.

In other words, Bill will be unhappy no matter what: The antecedent and consequent of (35) are irrelevant to each other.

4.5 | Interim summary

In light of the observations in this section regarding discourse particles and coherence relations in various types of texts, there is good reason to think that coherence establishment can account for the relevance effect where it occurs, and its non-occurrence in irrelevance conditionals. This conclusion is compatible with every non-inferentialist theory of conditionals that I am aware of. Furthermore, it undermines the major arguments for the inferentialist treatment of relevance as a conventional component of conditional meaning: Relevance and irrelevance inferences in conditionals can be explained by independent factors, with no need to postulate conventionalized relevance inferences or an ambiguity in conditionals.

The next sections address a number of further questions that arise in accounting for the range of arguments that have been produced in favor of inferentialism.

5 | WHY ARE MISSING-LINK CONDITIONALS ODD?

As noted above, one of the major arguments for inferentialism is the oddness of missing-link conditionals, which suggests that the relevance effect is not due to Gricean pragmatics. After all, most types of conversational implicature can be suspended when the context does not support them, but failures of relevance generate a sense of inappropriateness instead. ("??" indicates that an example is odd or infelicitous.)

38. ?? If Paris is the capital of France, Montreal is not in Chile.

39. If Mary is here, Bill is too. ?? But these things have nothing to do with each other.

However, this observation does not undercut a coherence-based theory, since coherence establishment is an obligatory part of interpretation. The literature on discourse coherence contains discussion of many examples with a similar flavor, where a text is felt to be infelicitous because it is difficult to infer a plausible coherence relation among clauses. For instance:

40. John broke his leg. I like plums.   (Knott & Dale, 1994)

This text is strikingly odd, but there is no hope of pinning its oddness on a conventional inference associated with some connective or other linguistic item: it is simply the juxtaposition of two sentences. The same holds of various other ways of connecting the sentence, including *if*.

41. a. John broke his leg, and/but/similarly I like plums.
    b. If John broke his leg, I like plums.

All of these oddities can be explained in the same way: We have a limited repertoire of coherence relations available when interpreting texts, and each places requirements on the relation between these two sentences that are implausible given typical background assumptions.

On the coherence theory, the contrast between the felicitous conditional (42a) (= (1)) and the missing-link conditional (42b) (= (4))—which we used above to motivate the relevance effect—is explained in the same way as the contrast in (43).

42. a. If Mary left the party early, Bill was unhappy.
    b. If Mary left the party early, Bill is Swedish.

43. a. Mary left the party early. Bill was unhappy.
    b. Mary left the party early. Bill is Swedish.

It is natural to read both (42a) and (43a) as suggesting that Bill was unhappy *as a result* of Mary's departure—a result relation—or else that Mary left *because* Bill was unhappy—an *explanation* relation. Neither of these possibilities is explicitly signaled by linguistic means; instead, we infer a plausible coherence relation between clauses as part of the process of rationalizing why a speaker would choose to connect these two clauses in constructing a text.

By contrast, (42b) and (43b) are strange because we are unable to infer a plausible coherence relation. An explanation relation is implausible given the content of the clauses, and a result relation is practically impossible—Mary's departure could hardly affect Bill's nationality. However, the coherence of these examples improves dramatically if we supply the background information that Mary has an irrational dislike of Swedes, making it plausible that *Bill is Swedish* does in fact stand in an explanation relation to *Mary left the party early*. Strikingly, the same background information is sufficient to render felicitous both the conditional (42b) and the juxtaposition of sentences in (43b). This suggests that the phenomenon of missing-link conditionals is reducible to a more general phenomenon of missing-link texts.

6 | THE DEFAULT STATUS OF THE RELEVANCE EFFECT

Why are "normal" conditionals associated with the relevance effect, while irrelevance conditionals are felt as "abnormal"? As Douven (2008) notes in the passage quoted in Section 3, the sense of normality invoked here is statistical: the claim is that a majority of conditionals in actual use enforce a relevance requirement, and that the inferentialist theory explains this.

We should be hesitant to adopt such a statistical claim without evidence from analysis of linguistic corpora. Still, the claim may be true. If it is, the inferentialist theory does *not* in fact account for it. Instead, this theory treats *if* as ambiguous between a reading that enforces relevance and one that does not, and gives no account of how or why one reading is chosen. The latter component is what would be needed to explain the purported statistical prevalence of relevance-enforcing readings.

In contrast, the coherence theory goes some way toward explaining why there should be a preference for relevance-enforcing interpretations (result and explanation) in conditionals that are not explicitly marked as violated expectation/denial of preventer. As we saw in Section 4.4, several different linguistic contexts seem to disfavor violated expectation/denial of preventer relations except when there is explicit marking to this effect. In conjunctions, conditionals, and bare juxtapositions of sentences, a result or explanation relation was preferred *even when* such a relation was implausible in light of the content. Here are some more examples:

44. a. John called Bill a nice guy. Bill punched him.
    b. John called Bill a nice guy, and Bill punched him.
    c. After John called Bill a nice guy, Bill punched him.
    d. Bill punched John. John had called him a nice guy.
    e. Bill punched John if John called him a nice guy.

All of these texts are naturally read as indicating that Bill punched John *because* John called him a nice guy. This is striking, because this causal relationship is inconsistent with our normal assumptions about typical human behavior. There seems to be a general default for inferring such causal relations rather than violated explanation/denial of preventer, even when there is no incoherence in the latter, and world knowledge supports them more strongly than result or expectation relations. The missing interpretations do, of course, come out when explicitly marked:

45. a. <u>Even though</u> John called Bill a nice guy, Bill punched him.
    b. If John called Bill a nice guy, Bill <u>still</u> punched him.

A general preference for result and explanation relations over violated expectation and denial of preventer goes some way toward accounting for the preference for relevance readings of conditionals, then: Result and explanation are incompatible with irrelevance readings. This style of explanation does not, however, account for why this preference would exist. Instead, it simply shows that the problem may be reducible to a larger unsolved problem.

While a careful study is beyond the scope of this article, I speculate that it is related to the content of the coherence relations themselves. If our expectations are well-attuned to the statistics of the real world, there *should* be more expectation-verifying than expectation-falsifying events to talk about. It may also be that the preference for result and explanation, absent explicit marking to the contrary, is particularly strong in conditionals because of the need to rationalize why a speaker would have chosen to subordinate the material in the consequent to the information contained in the antecedent, as opposed to using an informationally symmetric device such as *and* or juxtaposition. But these are merely speculations at this point, and further empirical and theoretical investigation is needed in order to verify that the posited preference is indeed real and that it extends to a variety of conditional and non-conditional contexts.

7 | AN EXPERIMENTAL STUDY ON COHERENCE AND RELEVANCE

The coherence-based account of the relevance effect was anticipated by Cruz et al. (2016), who suggested that missing-link conditionals are odd because of the lack of a "common topic of discourse." In the version proposed here, the problem in missing-link conditionals is not merely that the antecedent and consequent are about a different "topic" in some generalized sense. Instead, the establishment of discourse coherence relations relies on a certain inventory of available relations among clauses, a few of which were discussed in detail above.

Krzyżanowska et al. (2017) followed up on the account of Cruz et al., comparing conditionals to two-person conversational exchanges in a task designed to dissociate relevance from presence or absence of a common topic. In one condition they asked participants to judge to what extent it makes sense to give a certain response to a context sentence in three conditions. An example stimulus is given in (46).

46. [Patrick plans to take his girlfriend, Sophie, for short holidays. He is discussing different ideas with Matt. They have been to the Alps plenty of times, so now Patrick considers a hiking trip in the Pyrenees.]
    Patrick: "Sophie likes the Alps."
    a. Matt: "She will enjoy hiking in the Pyrenees."   (Positive relevance-same topic)
    b. Matt: "Mountaineering can be dangerous."   (Irrelevant-same topic)
    c. Matt: "More and more people in Western Europe care about animal welfare."
       (Irrelevant-different topic)

Krzyżanowska et al. (2017) also asked for "makes sense" ratings for matched conditionals in the same three conditions.

47. [Same context, but no assertion from Patrick.]
    a. Matt: "If Sophie likes the Alps, then she will enjoy hiking in the Pyrenees." (Positive relevance-same topic)
    b. Matt: "If Sophie likes the Alps, then mountaineering can be dangerous." (Irrelevant-same topic)
    c. Matt: "If Sophie likes the Alps, then more and more people in Western Europe care about animal welfare." (Irrelevant-different topic)

Krzyżanowska et al. (2017) reasoned that, if the relevance effect was due to the lack of a common topic of discourse as per Cruz et al. (2016), the conditionals in (47) should be affected in the same way as the dialogs in (46). If conditionals carry an additional relevance requirement, then the (a) response should make more sense than the (b) response in (47), with no such contrast in (46). They found precisely this difference, and interpreted it as evidence for inferentialism and against a coherence-based account.

There are two problems with this argument. The first is that the conditionals in (47) contain the item "then," which independently enforces a result or explanation relation—but the stimuli in (46) do not. However, it is likely that qualitatively similar results would appear if this confound were removed. More crucially, the problem is that coherence establishment is much more fine-grained than mere "common topichood," and it is constrained by semantic characteristics of the material used in a way that disrupts the parallel between conditionals and dialogs that Krzyżanowska et al.'s argument assumes.

An unmarked conversational exchange is compatible with a wide variety of coherence relations, for instance *Parallel*.

> **Parallel**: Infer $p(a_1, a_2, ...)$ from an assertion of S_1 and $p(b_1, b_2, ...)$ from an assertion of S_2, where for some property vector $\vec{q}$, $q_i(a_i)$ and $q_i(b_i)$ for all *i*. (Kehler, 2002, p. 16).

A different stimulus from the crucial *irrelevant-same* topic condition of (Krzyżanowska et al., 2017) provides a clear example:

48. [Laura and Kate discuss local politics. They are particularly worried about the city's government not investing in public transport.]
    Kate: "Public transport is inefficient."
    Laura: "There are hardly any bicycle lanes."

Here the parallel properties are *being inefficient* and *hardly existing*, both of which are instances of the property *p* of being a negative transport-related political issue. The parallel arguments are $a_1 = public\ transport$ and $b_1 = bicycle\ lanes$, which share the common property q_1 of being crucial to efficient transport.

Krzyżanowska et al. (2017) assume that a coherence-based non-inferentialist theory would render a conditional variant of (48) felicitous. However, this reasoning goes through only if the operative parallel relation is available with conditionals in general. If it is not, then the unavailability of parallel would be sufficient to explain the infelicity.

In fact parallel does appear to be unavailable in conditionals. For example, consider (49), which strongly favors a parallel relation.

49. Kamala likes Stanley Kubrick. Joe likes Francis Ford Coppola.

Here $a_1 =$ Kamala, $b_1 =$ Joe, and the common property q_1 is something like *prominent member of the current administration*; a_2 is Kubrick, b_2 is Coppola, and q_2 is *being a famous director*; and the relation that binds them together is $p = like$.

There is a striking contrast between these felicitous instances of parallel and a modified version involving a conditional.

50. ?? If Kamala likes Stanley Kubrick, Joe likes Francis Ford Coppola.

The only way to make sense of this example is to suppose that knowledge of Kamala′s feelings toward Kubrick somehow allows us to infer Joe's feelings about Coppola. This indicates that parallel has been excluded, even though it is available in the matched juxtaposition in (49). The observation that parallel relations are systematically allowed in juxtapositions but excluded from conditionals is enough to explain why participants found (48) more sensible than its conditional variant: The former can be given a coherent interpretation, while the latter cannot.

Most of the critical materials in the experiment of Krzyżanowska et al. (2017) yield to a similar analysis, though not all of them instantiate the parallel relation. Some, such as (46b), are instances of Kehler's *contrast* relation, which also appears to be incompatible with conditionals.

Why are conditionals incompatible with certain relevance relations that are available in conjunctions, monologs, and dialogs? Perhaps it is because inferentialism is correct: After all, this would have the observed effect of restricting the available relations to result and explanation. However, this explanation would not account for the additional availability of violated explanation and denial of preventer readings in conditionals, as we considered in some detail above. Furthermore, there is reason to suspect again that we are dealing with a very general discourse-related phenomenon, rather than an arbitrary lexical property of the item *if*. Specifically, other subordinating devices with similar meaning to *if* display *exactly the same restriction* as far as Parallel relations are concerned:

51. a. ?? Assuming/Provided Kamala likes Stanley Kubrick, Joe likes Francis Ford Coppola.
    b. ?? On the assumption that Kamala likes Stanley Kubrick, Joe likes Francis Ford Coppola.
    c. ?? Supposing Kamala likes Stanley Kubrick, Joe likes Francis Ford Coppola.

With these and a variety of other devices that flag a piece of information as a temporary assumption, we find that the examples make sense only if there is an inferential connection between the politicans′ feelings about the various directors—that is, if we read the text as involving explanation or result. (As the reader can verify, the addition of a particle such as *still* makes available a violated explanation reading with all of these suppositional devices—which, however, remains infelicitous into the examples under discussion.)

The shared meaning component between these items and *if* involves the discourse attitude of *supposition*—evaluation of the consequent in a context where we have temporarily assumed the truth of the antecedent. In contrast, it is possible in dialogs, monologs, and conjunctions to treat adjacent clauses as independent pieces of information, neither subordinated to the other. This independence is required by the parallel relation, simply in terms of its meaning.

This discussion only scratches the surface of the rich interaction between coherence relations and coordinating/subordinating devices, and it remains to be seen what further restrictions are associated with conditional meaning. However, it is clear that conditionals and other suppositional devices impose restrictions on available coherence relations that are not present in dialogs, and the generality of this effect suggests that it is not simply an arbitrary lexical property of *if*. This means that the greater acceptability of certain dialogs vs. matched conditionals does not, in itself, present an argument against a coherence-based explanation of the relevance effect.

8 | STALNAKER'S THESIS AND STRONG CENTERING

As noted in Section 1, inferentialism problematizes two important properties shared by many theories of conditionals: Stalnaker's thesis and strong centering. Crucially, neither property makes reference to relevance between A and C.

Stalnaker's thesis: $P(\text{If A, C}) = P(\text{C}|\text{A})$.

Strong centering: If A and C are both true, then "If A, C" is true.

There is a great deal of empirical evidence supporting Stalnaker's thesis (Douven & Verbrugge, 2010, 2013; Fugard et al., 2010; Hadjichristidis et al., 2001; Over & Evans, 2003). However, if inferentalism is correct, it should be easy to generate counter-examples: Simply find cases where the probability of C is high, and A and C are probabilistically independent, and so irrelevant to each other. Then P(C|A) should be high—equaling P(C)—but P(If A, C) should be low because of the lack of connection between antecedent and consequent. For instance, (52) should have low probability even though the conditional probability of the consequent given the antecedent is high.

52. If John broke his leg, I am not asleep.

Skovgaard-Olsen et al. (2016) manipulated probability and relevance in this way, finding that participants' judgments of the probability of conditionals were generally lower when the antecedent was irrelevant to the consequent, and even lower when it was negatively relevant (that is, decreased the probability of the consequent compared with its unconditional probability). They interpreted this finding as evidence that Stalnaker's thesis holds only in cases of positive relevance. This interpretation is compatible with inferentialism, but not with Stalnaker's thesis in general.

Similarly, strong centering is highly plausible and enjoys some empirical support (Cruz et al., 2015; Politzer & Baratgin, 2016). However, inferentialism implies that there should be counter-examples whenever A and C both happen to be true, but they have nothing to do with each other. If John did break his leg, and I am not asleep, then inferentialism predicts that (52) is false even though both antecedent and consequent are true. In addition, the conjunction of antecedent and consequent may have high probability even though the conditional has low probability. This violates the inequality P(If A, C) ≥ P(A and C) that is enforced by classical probability together with strong centering. When Skovgaard-Olsen, Kellen, et al. (2019) tested these predictions experimentally, they found that participants were indeed less confident in the

inference from the *and* sentence to the *if* sentence under irrelevance, and that participants sometimes assigned higher probability to "A and C" than to the conditional "If A, C."

While both of these sets of results are compatible with inferentialism, they are equally compatible with the coherence-based theory proposed here. The crucial fact is that there is considerable independent evidence that judgments of truth and probability in experimental settings are not "pure" in the sense of depending exclusively on conventional, literal meaning as theorized in natural language semantics. Instead, participant judgments frequently take into account additional, non-literal aspects of conveyed meaning such as conversational implicatures. The experimental pragmatics literature is replete with examples. One of the most common ways to probe implicature computation is to observe when participants judge a sentence "false" ("wrong," "incorrect," etc.) because it is associated with a false implicature, even though its literal meaning is true (De Neys & Schaeken, 2007; Degen & Tanenhaus, 2015; Noveck, 2001, among many others). For instance, Bott and Noveck (2004, Experiment 3) found that only 41% of participants agreed that "Some elephants are mammals" is true, compared with 88% for "Some mammals are elephants." Even though the implicature associated with "Some elephants are mammals" is in principle optional, many participants perceived it nonetheless—perhaps because of the obvious true competitor with "all"—and included it in their judgments of truth and falsity.

In general, experimental participants' understanding of terms like "true" and "false" differs dramatically from the way that these terms are used in logic and semantics. (See Jasbi et al., 2019 for a recent discussion of the complexity of interpreting experimental measures in experimental semantics/pragmatics.) They appear to target the intended message of an utterance, including rich contextual and discursive information. Given this, the coherence-based theory of the relevance effect has an obvious strategy for dealing with apparent failures of strong centering. Just as participants sometimes judge a sentence false when its literal content is true but it is associated with a false implicature, they may judge a conditional false if its appropriate use in context would require accommodating a coherence relation whose preconditions are false. A strong prior belief that A and C are irrelevant, or negatively relevant, is incompatible with the preconditions of the result and explanation relations. But, as we saw in Section 6, there is a strong default toward exactly these coherence relations in conditionals that are not explicitly marked otherwise. Given this, it is not surprising that many participants sometimes judge a conditional false when it is literally true, since they default to a coherence relation whose background conditions are believed to be false.

A similar account is available for the apparent failures of Stalnaker's thesis that Skovgaard-Olsen et al. (2016) use to argue that relevance is part of the literal meaning of conditionals. If pragmatic considerations intrude on probability judgments as they do on truth-value judgments, a similar story can be given about cases where context and prior expectation do not support the inferred coherence relations. In fact, there is evidence that probability judgments are indeed sensitive to non-literal meaning. When Cremers et al. (2017) investigated probability assignments to sentences with false scalar implicatures, they found that about half of participants reported the probability of the sentence without the implicature, while another half reported the probability of the conjunction of the sentence and its implicature. This is what we would expect if participants vary in whether they compute an implicature, as in Bott & Noveck, 2004, and then judge the probability of the resulting (literal or pragmatically enriched) interpretation. Along similar lines, a coherence-based account of the relevance effect could posit that some participants assign low probability to a conditional if they interpret it as invoking a result or explanation relation, and they assign low probability

to the conjunction of the conditional's meaning and the background conditions of the inferred coherence relation.

Note that nothing in this section constitutes an argument *for* either strong centering or Stalnaker's thesis. The point is merely to show that the coherence-based explanation of the relevance effect makes available new explanations for experimental results that have been claimed to falsify these conditions. The primary arguments for adopting these conditions involve independent theoretical and empirical considerations.

9 | CONCLUSION

Many authors have recently argued that the frequent requirement of relevance between a conditional's antecedent and its consequent cannot be explained by Gricean pragmatics. This conclusion may be sound, but we cannot conclude that relevance is a conventional aspect of the meaning of conditionals. An overlooked third possibility, discourse coherence, turns out to provide a unified explanation of where the relevance requirement occurs and where it does not, with support from the distribution of various connectives and discourse particles. The crucial observation is that similar relevance and irrelevance inferences arise in non-conditional sentences, with similar options for support from connectives and discourse particles. Missing-link conditionals are a special case of the broader phenomenon of incoherent texts: they are conditionals in which it is not possible to supply a plausible coherence relation.

This decomposition of the relevance effect into conditional semantics and discourse coherence overturns the major arguments for inferentialism. It is, however, compatible with theories of conditionals that do not encode any information about relevance. This conclusion also suggests a response to arguments against important features of conditionals—strong centering and Stalnaker's thesis—that are otherwise theoretically and empirically well-motivated.

DATA AVAILABILITY STATEMENT
There is no data available.

ORCID
Daniel Lassiter https://orcid.org/0000-0001-9290-4274

REFERENCES
1. Adams, E. W. (1975). *The logic of conditionals: An application of probability to deductive logic*. Springer.
2. Asher, N. (1993). *Reference to abstract objects in discourse*. Kluwer.
3. Asher, N., & Lascarides, A. (2003). *Logics of conversation*. Cambridge University Press.
4. Bennett, J. F. (2003). *A philosophical guide to conditionals*. Oxford University Press.
5. Biezma, M. (2014). The grammar of discourse: The case of then. In T. Snider, S. D'Antonio, & M. Wiegand (Eds.), *Semantics and linguistic theory (SALT)* (Vol. 24, pp. 373–394). LSA and CLC Publications.
6. Bott, L., & Noveck, I. A. (2004). Some utterances are underinformative: The onset and time course of scalar inferences. *Journal of Memory and Language*, *51*(3), 437–457.

7. Chafe, W. (1987). Cognitive constraints on information flow. In R. S. Tomlin (Ed.), *Coherence and grounding in discourse* (Vol. 11, pp. 21–51). John Benjamins.

8. Chierchia, G., Fox, D., & Spector, B. (2012). The grammatical view of scalar implicatures and the relationship between semantics and pragmatics. In K. V. Heusinger, C. Maienborn, & P. Portner (Eds.), *Semantics: An international handbook of natural language meaning* (Vol. 3, pp. 2297–2232). Berlin: Walter de Gruyter.

9. Clark, H. H. (1996). *Using language*. Cambridge University Press.

10. Cremers, A., Križ, M., & Chemla, E. (2017). Probability judgments of gappy sentences. In S. Pistoia-Reda & F. Domaneschi (Eds.), *Linguistic and psycholinguistic approaches on implicatures and presuppositions* (pp. 111–150). Palgrave Macmillan.

11. Cruz, N., Baratgin, J., Oaksford, M., & Over, D. E. (2015). Bayesian reasoning with ifs and ands and ors. *Frontiers in Psychology*, *6*, 192.

12. Cruz, N., Over, D., Oaksford, M., & Baratgin, J. (2016). Centering and the meaning of conditionals. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the thirty-seventh annual conference of the cognitive science society*. Cognitive Science Society.

13. Davis, W. A. (1983). Weak and strong conditionals. *Pacific Philosophical Quarterly*, *64*(1), 57–71.

14. De Neys, W., & Schaeken, W. (2007). When people are more logical under cognitive load: Dual task impact on scalar implicature. *Experimental Psychology*, *54*(2), 128–133.

15. Degen, J., & Tanenhaus, M. K. (2015). Processing scalar implicature: A constraint-based approach. *Cognitive Science*, *39*(4), 667–710.

16. Douven, I. (2008). The evidential support theory of conditionals. *Synthese*, *164*(1), 19–44.

17. Douven, I. (2015). *The epistemology of indicative conditionals: Formal and empirical approaches*. Cambridge University Press.

18. Douven, I. (2016). On de Finetti on iterated conditionals. In C. Beierle, G. Brewka, & M. Thimm (Eds.), *Computational models of rationality: Essays dedicated to Gabriele Kern-Isberner on the occasion of her 60th birthday* (pp. 265–279). College Publications.

19. Douven, I. (2017). How to account for the oddness of missing-link conditionals. *Synthese*, *194*(5), 1541–1554.

20. Douven, I., Elqayam, S., Singmann, H., & van Wijnbergen-Huitink, J. (2018). Conditionals and inferential connections: A hypothetical inferential theory. *Cognitive Psychology*, *101*, 50–81.

21. Douven, I., & Verbrugge, S. (2010). The Adams family. *Cognition*, *117*(3), 302–318.

22. Douven, I., & Verbrugge, S. (2013). The probabilities of conditionals revisited. *Cognitive Science*, *37*(4), 711–730.

23. Edgington, D. (1995). On conditionals. *Mind*, *104*(414), 235–329.

24. Frank, M., & Goodman, N. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998.

25. Fugard, A. J., Pfeifer, N., Mayerhofer, B., & Kleiter, G. D. (2010). How people interpret conditionals: Shifts towards the conditional event. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(3), 635–648.

26. Grice, H. P. (1989a). Indicative conditionals. In P. Grice (Ed.), *Studies in the ways of words*. Harvard University Press.

27. Grice, H. P. (1989b). Logic and conversation. In P. Grice (Ed.), *Studies in the ways of words*. Harvard University Press.

28. Grice, H. P. (1989c). Further notes on logic and conversation. In P. Grice (Ed.), *Studies in the way of words*. Harvard University Press.

29. Hadjichristidis, C., Stevenson, R. J., Over, D. E., Sloman, S. A., Evans, J., & Feeney, A. (2001). On the evaluation of "if p then q" conditionals. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the twenty-third annual conference of the cognitive science society* (pp. 409–414). Lawrence Erlbaum Associates.

30. Hobbs, J. R. (1979). Coherence and coreference. *Cognitive Science*, *3*(1), 67–90.

31. Hobbs, J. R. (1990). *Literature and cognition*. CSLI Publications.

32. Iatridou, S. (1993). On the contribution of conditional *then*. *Natural Language Semantics*, *2*(3), 171–199.

33. Jasbi, M., Waldon, B., & Degen, J. (2019). Linking hypothesis and number of response options modulate inferred scalar implicature rate. *Frontiers in Psychology*, *10*, 189.

34. Kehler, A. (2002). *Coherence, reference, and the theory of grammar*. CSLI Publications.

35. Knott, A., & Dale, R. (1994). Using linguistic phenomena to motivate a set of coherence relations. *Discourse Processes*, *18*(1), 35–62.

36. Knott, A., & Dale, R. (1996). Choosing a set of coherence relations for text generation: A data-driven approach. In G. Adorni & M. Zock (Eds.), *Trends in natural language generation: An artificial intelligence perspective* (pp. 47–67). Springer.

37. Krzyżanowska, K. (2019). What is wrong with false-link conditionals? *Linguistics Vanguard*, *5*(3), 1–9.

38. Krzyzanowska, K., Collins, P., & Hahn, U. (2017). The puzzle of conditionals with true clauses: Against the Gricean account. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. J. Davelaar (Eds.), *Proceedings of the thirty-ninth annual meeting of the cognitive science society* (pp. 2476–2481). Cognitive Science Society.

39. Krzyżanowska, K., Collins, P. J., & Hahn, U. (2017). Between a conditional's antecedent and its consequent: Discourse coherence vs. probabilistic relevance. *Cognition*, *164*, 199–205.

40. Krzyżanowska, K., & Douven, I. (2018). Missing-link conditionals: Pragmatically infelicitous or semantically defective? *Intercultural Pragmatics*, *15*(2), 191–211.

41. Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, *85*(3), 297–315.

42. Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*(1), 339–359.

43. Lycan, W. G. (2001). *Real conditionals*. Clarendon Press.

44. Noveck, I. A. (2001). When children are more logical than adults: Experimental investigations of scalar implicature. *Cognition*, *78*(2), 165–188.

45. Over, D. E., & Evans, J. S. B. (2003). The probability of conditionals: The psychological evidence. *Mind & Language*, *18*(4), 340–358.

46. Politzer, G., & Baratgin, J. (2016). Deductive schemas with uncertain premises using qualitative probability expressions. *Thinking & Reasoning*, *22*(1), 78–98.

47. Quine, W. V. O. (1965). *Elementary logic*. Harper.

48. Schiffrin, D. (1987). *Discourse markers*. Cambridge University Press.

49. Schiffrin, D. (2001). Discourse markers: Language, meaning, and context. In D. Schiffrin (Ed.), *The handbook of discourse analysis* (pp. 54–75). Wiley.

50. Skovgaard-Olsen, N. (2016). Motivating the relevance approach to conditionals. *Mind & Language*, *31*(5), 555–579.

51. Skovgaard-Olsen, N., Collins, P., Krzyżanowska, K., Hahn, U., & Klauer, K. C. (2019). Cancellation, negation, and rejection. *Cognitive Psychology*, *108*, 42–71.

52. Skovgaard-Olsen, N., Kellen, D., Hahn, U., & Klauer, K. C. (2019). Norm conflicts and conditionals. *Psychological Review*, *126*(5), 611.

53. Skovgaard-Olsen, N., Singmann, H., & Klauer, K. C. (2016). The relevance effect and conditionals. *Cognition*, *150*, 26–36.

54. Stalnaker, R. C. (1970). Probability and conditionals. *Philosophy of Science*, *37*(1), 64–80.

55. van Rooij, R., & Schulz, K. (2019). Conditionals, causality and conditional probability. *Journal of Logic, Language and Information*, *28*(1), 55–71.