Conditional answers and the role of probabilistic epistemic representations

Jos Tellings
Utrecht University / Utrecht Institute of Linguistics
Trans 10, 3512 JK Utrecht
the Netherlands
`j.l.tellings@uu.nl`

Abstract

Conditional utterances can be used in discourse as answers to regular, non-conditional questions in situations of partial knowledge of the answerer. I claim that the probabilities that interlocutors assign to each other's possible epistemic states are a measure of the relevance of conditional answers. A second criterion that makes a conditional answer 'if p, then q' relevant has to do with the dependency between p and q that is conveyed in the statement. A conditional answer counts as relevant when this dependency leads the question asker to shift from a decision problem about q to an alternative, easier, decision problem about p.

1 Introduction

1.1 Conditionals as answers

The study of conditional sentences (*if-then* sentences) constitutes a vast area within formal semantics. A lot of this work considers the internal structure of conditionals: how do tense, aspect, and modality contribute to the combination of pragmatic and semantic effects observed for conditional utterances? (see e.g. von Fintel, 2011 for an overview). The use of conditionals in conversation is relatively less well studied. Yet, corpus research shows that conditionals are a very common utterance type (Ferguson, 2001; for some preliminary corpus data from the Europarl corpus, see Tellings, 2020). A prominent view in semantic theories of discourse interaction is to model conversational contexts as a structured collection of (possibly implicit) information requests, and answers to these requests. These requests are called *questions under discussion* (QUDs, Roberts, 1996/2012), or *issues* in the framework of Inquisitive Semantics (Ciardelli et al., 2018). Hence, a natural start of a theoretically and computationally robust study of condi-

tionals in discourse is to ask how conditional utterances behave as **answers to questions**.

Work on conditionals as answers to questions is surprisingly scarce, given that both question-answer models of discourse, and formal theories of the meaning of conditional sentences, have been major themes in the literature on semantics and pragmatics. I argue that the most important empirical setting to study is when conditionals answer regular, non-conditional questions. Indeed, we observe that conditional utterances make good answers to all types of non-conditional questions, including polar questions (1a), alternative questions (1b) and *wh*-questions (1c):

(1) a. A: Will John come to the party?
        B: If he finishes his work, he will.

    b. A: Do you want coffee or tea?
        B: If it is freshly made, I would like coffee.

    c. A: What will John cook for dinner?
        B: If he managed to buy parmesan cheese, he will make pasta.

Here, the conditional form of the answer was not driven by the form of the question, but newly introduced by B on the basis of conditional knowledge he has.[1] What is worth noting about the cases in (1), is that B's answers are not the maximally informative *congruent* answers that are typically studied in semantic theories of questions (congruent answers would be 'yes'/'no' to a polar question, 'coffee' or 'tea' in (1b), and a constituent answer such as 'pasta' or 'steak' in (1c)). This is because the conversations in (1) crucially involve *partial knowledge*. In (1a), B does not know for

---

[1]Throughout, I will use A(lice) and female pronouns to refer to the question asker, and B(ob) and male pronouns to refer to the answerer.

a fact whether John will come, so the pragmatically most informative answers 'Yes' and 'No' are ruled out. However, he is not completely ignorant about the situation, he has *conditional knowledge*: B knows that if John finishes his work, he will come.

From these observations, my central research question emerges: in what situations do speakers give a conditional answer, and how do speakers choose a particular conditional answer as their optimal response in comparison to other conditional and non-conditional answer options?

These questions can be made more explicit by introducing some theoretical concepts on the relevance of answers. In game-theoretic models of Gricean pragmatics (see Benz and Stevens, 2018 for an overview), the relevance of an answer with respect to a given question has been formalized by modeling the question as a decision problem, and assigning potential answers a **utility value**. Different variants of game-theoretic pragmatics (e.g. older utility-based accounts (van Rooij, 2004), optimal answer models (Benz and van Rooij, 2007), and rational speech act (RSA) models (Goodman and Frank, 2016)) use the notion of utility in different ways, but the common view is that B chooses an answer that maximizes utility with respect to the decision problem A is trying to resolve.

Framed in these terms, the aim is to assign a utility value to conditional statements. By considering some specific examples, I will propose two probability measures that are proportional to the utility of a conditional answer. The first, discussed in section 2, is an epistemic condition related to the probabilities assigned to possible epistemic states of the interlocutors. The second, discussed in section 3, has to do with the dependency between the propositions expressed in *if*-clause and main clause. There we will also see that the simple approach of applying an existing utility-based framework to a material conditional proposition of the form '$p \supset q$' does not work.

Before all this, we will look at conditional answers from a somewhat different domain, conditional perfection (in §1.2), and show that two approaches to conditional answers in earlier literature do not meet the desiderata that I have outlined above (in §1.3).

1.2 Conditional perfection

An additional reason to study conditional answers comes from the pragmatic phenomenon of **conditional perfection**, or the strengthening of conditionals to biconditionals. The following datum from van Canegem-Ardijns and van Belle (2008) is an example:

(2) If you pay your contribution, you may participate in the barbecue.
implicature: if you don't pay, you may not participate

Conditional perfection is widely studied in the pragmatics literature (Geis and Zwicky, 1971; de Cornulier, 1983; von Fintel, 2001; van Canegem-Ardijns, 2010; among others), and various different mechanisms for deriving the implicature have been proposed (see e.g. van der Auwera, 1997 for an overview). A major question in the work on conditional perfection is why the implicature arises in some cases, such as (2), but not in others, such as (3) (from von Fintel, 2001):

(3) If this cactus grows native to Idaho, then it is not an *Astrophytum*.
$\nrightarrow$ If this cactus doesn't grow native to Idaho, it is an *Astrophytum*.

In more recent work on conditional perfection, it has been proposed that perfection occurs when a conditional is interpreted as an exhaustive answer to the question under discussion (Herburger, 2015, cf. von Fintel, 2001; see Cariani and Rips, 2017 for an experimental approach to this idea).

(4) If you work hard you will succeed.
Exhaustification: ⟨. . . and only if you work hard you will succeed⟩
(Herburger, 2015)

In unrelated work on exhaustive answers, it has been proposed that whether an answer is interpreted as exhaustive or not (*mention-all* or *mention-some*) depends on "human concerns" underlying the asking of the question, which can again be modeled in game-theoretic pragmatical models in terms of the decision problem the speaker is trying to solve (van Rooij, 2004).

Hence, in order to understand conditional perfection better, we need to understand when conditionals are interpreted as mention-some answers,

and how conditional answers correspond to the speakers' interests. Therefore, a study of the utility of conditional answers with respect to the interlocutors' interests can contribute to the understanding of conditional perfection.

1.3 Earlier work

In order to further appreciate the approach to conditionals taken here – as answers to regular questions –, it is worth briefly reviewing earlier work on conditional answers. I will mention two lines of work, which however in my opinion are not representative of the wider problem of conditional sentences in discourse that I address here.

The first is work on conditional utterances as answers to conditional questions, as in example (5) from Isaacs and Rawlins (2008, 276):

(5) A: If Alfonso comes to the party, will Joanna leave?
     B: If he comes, Joanna will leave.

The reason that these types of question-answer combinations are not very insightful for my purposes, is that here the conditional answer merely mimics the conditional form of the question. Hence, B will have had no independent grounds to choose a conditional form for his answer.

In the same vein, Ippolito (2013) proposes that counterfactual conditionals are answers to conditional questions under discussion (CQUDs), as the following example taken from her paper illustrates:

(6) [CQUD: If the weather had been fine, would Jones be wearing his hat?]
     If the weather had been fine, Jones would be wearing his hat.

This illustrates my point that studying conditionals as answers leads to a better overall understanding of conditional statements in general, because these CQUDs are generally implicit, and Ippolito's work is not part of understanding (counterfactual) conditionals in discourse. However, the same point about form parallelism in question and answer can be made for (6). Moreover, Ippolito (2013) does not take into consideration that the *if*-clause and main clause of a conditional can have different information structural statuses, depending on how they are used in a dialogue context. In fact, both *if*-clause and main clause may be in focus (see e.g. Farr, 2011; Tellings, 2016, §4.4).

The second line of work I want to mention here is Hesse et al. (2018), because it is methodologically closer to what I aim to do (cf. also Stevens et al., 2016), but studies a different kind of conditional expression, namely *speech act conditionals*/SACs (also known as *biscuit conditionals*). They give the following example, illustrating 'positive', 'negative', and 'alternative' speech act conditionals as answers to a polar question:

(7) A: Is there a restaurant close to the apartment?

     a. B: If you enjoy eating out, there is an Italian restaurant in the street.   [PSAC]

     b. B: If you enjoy eating out, there is an Italian restaurant in the neighboring quarter.   [NSAC]

     c. B: If you enjoy eating out, there is an Italian restaurant as well as a food court nearby.   [ASAC]

Hesse et al. provide a model based on this specific example of a client asking a real estate agent questions about an apartment that predicts when a SAC is generated. The model does not, however, address the issue of why the answer is expressed by a SAC, and not by some other linguistic construction. For example, the response to the question in (7) could also be expressed as in (8B):

(8) A: Is there a restaurant close to the apartment?
     B: Ah, you like eating out? Yes, there is an Italian restaurant in the street.

Hesse et al. state that the choice between SACs in (7) and answers such as (8B) "depends on discourse-dependent and stylistic reasons" (p. 103), but do not elaborate.

In conclusion, both the work on conditional questions, and the work on speech act conditional answers, take a rather limited view on conditional answers, and does not take into account the general relationship between the two.

2   An epistemic licensing condition

I will start by looking at some specific situations in which conditional answers are licensed.

First, let me set out the boundary conditions as introduced in the previous section. I assume that A asks a question $?q$, and B answers 'if p, (then)

q', written as '$p \to q$'. It follows that $\neg K_A?q$ (this is the standard assumption of speaker ignorance: A does not ask a question if she knows the answer to it already). As pointed out above, conditional answers are uttered in situations of partial knowledge, so I moreover assume that $\neg K_B?q$ (otherwise, B would have given a complete answer 'yes' or 'no'). Finally, I assume that $\neg K_B?p$ (this is a standard presumption on conditional utterances; if B knows whether p, he wouldn't utter a conditional sentence with 'if p'). These conditions can be seen in the following example of a licensed conditional answer:

(9)   [Alice calls to the IT help desk]
     A: Did I install my printer correctly?
     B: If there is a printer icon on the desktop, you installed it correctly.

The three conditions introduced above are satisfied, because A doesn't know whether she installed her printer correctly, B doesn't know either (he is at a different place), and B doesn't know whether there is a printer icon on A's desktop (B doesn't have access to A's screen).

I claim that the felicity of B's answer in (9) has to do with the fact that B knows that A knows, or can easily verify, whether there is a printer icon on her desktop. More generally, I claim that the utility of a conditional answer depends on the probability that B assigns to that A knows about the antecedent p, more formally on $P_B(K_A?p)$:[2] the higher this value is, the more useful the conditional answer. The specific context in (9) represents the 'extreme' case in which $P_B(K_A?p) = 1$ (or sufficiently close to 1 by contextual standards), i.e. $K_B K_A?p$. Here $K_B K_A?p$ is a reasonable assumption, given that B knows that A has access to her screen and can easily verify the truth value of p.

In this epistemic setting, B entertains two possible candidates for A's epistemic state: one in which $K_A p$, and one in which $K_A \neg p$.[3] In the for-

mer case ($K_A p$), '$p \to q$' is a highly useful answer, because by modus ponens, A can conclude that q, and solve her decision problem $?q$. In the case that $K_A \neg p$, the answer '$p \to q$' is useful when it undergoes conditional perfection. In that case, A concludes that $\neg q$ by modus ponens. This links to the problem of characterizing the distribution of conditional perfection described in §1.2 above. This suggests the possibility of a second process that leads to conditional perfection: not only the exhaustive interpretation of a conditional answer (recall (4) above), but also a type of backwards reasoning on the part of A of the steps just described. Informally, A gets the perfection implicature, because she reasons that B's conditional answer is only relevant in the $\neg q$-state when perfection happens. Various details of this proposal about conditional perfection need to be worked out, which I leave aside for now.

Observe that a close variant of the conversation in (9) exists, in which instead of 'if p', an intermediate question $?p$ is uttered (B_1 below):

(10)   A: Did I install my printer correctly?
     B_1: Is there a printer icon on the desktop?

   a.   A: Yes.
       B_2: Then you installed it correctly.

   b.   A: No.
       B_2: Then there is a problem.[4]

This shows that there is a connection between conditionalizing an answer with 'if p' (in (9)), and asking an intermediate question $?p$ (in (10)). One of the default pragmatic conditions for a speaker A asking a question $?q$ to a speaker B, is ADDRESSEE COMPETENCE. This refers to the condition that A thinks that B is likely to know the answer, or in other words that $P_B(K_A?q)$ is high. If you want to know what the French word for 'rhubarb' is, you better ask somebody who knows French, rather than somebody who doesn't know French. Of course, there are exceptions to this default rule (e.g. in exam questions), but in general, this assumption is often satisfied: either trivially in addressee-directed questions such as 'How are you?', 'What did you do yesterday?', etc., or in the case of more factual questions as in (9). Literature on the pragmatics of questions has mostly ignored the condition of ADDRESSEE COMPETENCE, pre-

---

[2]This mixture of probability and epistemic logic is formalized in models such as van Benthem et al. (2009). The '$?p$' notation comes from inquisitive epistemic logic (Ciardelli and Roelofsen, 2015), in which the equivalence $K_A?p \leftrightarrow (K_A p \vee K_A \neg p)$ comes out as a logical validity, not an abbreviation.

[3]The relation between the epistemic representation '$K_A?p$' and the two candidates for epistemic states of A that B entertains, is formally present in inquisitive epistemic logic (see fn. 2), since $K_A?p$ abbreviates $K_A\{p, \neg p\}$. The same holds for other types of questions.

[4]Observe that in both (10a) and (10b), B_2 requires some sort of anaphoric expression 'then' or 'in that case'.

sumably because selecting which person you will ask your question to seems to be an issue that falls outside of linguistics proper. However, the data presented here show that the ADDRESSEE COMPETENCE condition ('$P_B(K_A?p)$ is high') is in fact linguistically relevant, because it doubles as a licensing condition for conditional answers with 'if p' in the given epistemic situation.

The context in (9), in which $P_B(K_A?p)$ takes its maximal value of 1, is perhaps a somewhat uncommon situation. It is instructive to consider cases toward the other end of the scale, where $P_B(K_A?p)$ is low, or indeed 0 (i.e. B knows for a fact that A does not know whether p). Two cases illustrating this epistemic situation are given below.

(11)   A: Did it rain yesterday?
        B: #If the atmospheric pressure was no higher than 1020 mBar and the squall line progression halted over Western Massachusetts, it did.

(12)   [epistemic situation: B knows that A is unaware of John's work situation]
        A: Will John come to the party?   (=(1a))
        B: If he finishes his work, he will.

Example (11) illustrates a situation in which a conditional answer is pragmatically odd (even though it may be true). Here B knows that A does not know whether the proposition in the *if*-clause is true, nor does A have an easy way to verify its truth or falsity (unless A is a professional meteorologist). Example (12), on the other hand, contains a conditional answer that may be licensed in the same epistemic situation without problem: imagine that B knows for a fact that A has no knowledge about John's work situation. This does not render the conditional answer unacceptable, despite the fact that the epistemic condition proposed in this section has not been satisfied.

I will come back to the difference between (11) and (12) in section 3, but to complete the line of argumentation, consider the following. That (12) is indeed licensed in the situation given above, and is different from the earlier example in (9), can be tested by using the connection between (9) and (10) as a diagnostic: for (12), changing the *if*-clause into an intermediate question does not work, see (13).

(13)   [epistemic situation: B knows that A is unaware of John's work situation]
        A: Will John come to the party?
      #B_1: Did he finish his work?
        ⋮

In the epistemic context just sketched, ADDRESSEE COMPETENCE is violated for B_1.[5] The fact that (12) is nonetheless an acceptable conditional answer shows that in addition to the epistemic condition proposed in this section, there is a second way in which conditional answers can be licensed. This second condition, however, does not license intermediate questions.

The broader conclusion of this section is that the discussed data offer insights into the sort of information that interlocutors keep track of in the course of conversation. In a theory of discourse dynamics such as Farkas and Bruce (2010), interlocutors keep track of each other's discourse commitments. I have shown that this should also include the representations of each other's epistemic states, here represented by probability distributions over their epistemic commitments. Speakers have prior representations about each other's epistemic states, including certain and presumed knowledge about other speakers' knowledge. These representations get updated over the course of the conversation, as new information is provided and new issues are raised. Interlocutors keep track of what other speakers know, and reason about this probabilistically. This is an aspect of probability in meaning that, to the best of my knowledge, has not been addressed in earlier literature.

3   Conditional dependency as relevance

I claim that when the epistemic licensing condition from section 2 is not fulfilled, as in the setting of example (12), a conditional answer can still be relevant, namely when the conditional dependency between p and q that is conveyed by the conditional utterance is relevant information.

The problem in formalizing this dependency is that the belief in a conditional statement 'if p, (then) q' is not simply equal to the conditional probability of q given p, as Lewis (1976) famously showed. This is also the reason why many game-theoretic pragmatic accounts, which are based on

---

[5]A similar observation can be made for (11): the intermediate question 'Was the atmospheric pressure [...] Western Massachusetts?' is unacceptable.

probabilities of utterances, cannot be straightforwardly applied to conditional answers.

Van Rooij and Schulz (2019) propose an assertability condition for conditional utterances that takes these problems into account:

(14) Van Rooij and Schulz's (2019) assertability condition for conditionals '$p \rightarrow q$':

$$\Delta^* P_p^q := \frac{P(q|p) - P(q|\neg p)}{1 - P(q|\neg p)} \quad \text{should be high.}$$

This condition incorporates the idea that a conditional statement should convey a dependency between p and q. For example, it rules out cases such as:

(15) #If it is sunny today, Jan Ullrich won the Tour de France in 1997.

(15) comes out as true by virtue of its consequent being true, but the conditional is not assertable because there is no (causal) dependency between *if*-clause and main clause (see van Rooij and Schulz, 2019 for details on the link between conditional probabilities and causality).

Van Rooij and Schulz use their condition (14) as a criterion for asserting conditionals, but I claim it is also used in the other direction: updating one's belief state upon hearing a conditional answer.

Utility of a conditional Can this general notion of relevance – conditional dependency as relevant information – be expressed in terms of the notion of utility?

In utility-based frameworks, one starts with a utility function $U(a, w)$ that assigns a utility value to an action a in world w. Then, the notion of the expected utility of an action a given a proposition f is introduced (Benz and van Rooij, 2007):

$$\text{EU}(a|f) = \sum_w P(w|f) \cdot U(a, w).$$

The next step is to define a notion of utility value of a message. This is done in different ways by different authors, but below is one proposal (Benz and van Rooij, 2007, 67):

$$\text{UV}(f) = \max_i \text{EU}(a_i|f).$$

I argue that one cannot simply represent the utility of a conditional answer as $\text{UV}(p \supset q)$, i.e. by

computing the utility value of the material conditional. I will identify three problems with this approach.

The first problem is that the truth conditions of the material conditional do not include the notion of conditional dependency as in (14) (the difference between truth conditions and assertability conditions is made clear by examples such as (15), see van Rooij and Schulz, 2019 for further discussion).

A second issue is that the conditional answers we have seen so far do not look like conditional propositions, but rather like conditional speech acts (or *conditional assertions*, see Stalnaker, 2011).[6] This intuition is supported by the observation that the consequent of conditional answers need not be a full proposition, but can take the form of a fragment answer, or in the case of the polar question (16a), 'yes' or 'no':

(16) a. A: Will John come to the party?
B: If he finishes his work, <u>yes</u>.

b. A: Do you want coffee or tea?
B: If it is freshly made, <u>coffee</u>.

c. A: What will John cook for dinner?
B: If he managed to buy parmesan cheese, <u>pasta</u>.

In each case, the consequent fragment answer (underlined in (16)) is, by itself, a valid answer to A's question as a whole. This suggests that the function of the *if*-clause is to conditionalize the answering speech act, rather than forming a conditional proposition $p \supset q$.

Finally, a more general problem with the decision problem approach is that the utility function and the set of actions is usually considered to be a fixed part of the model (a typical definition of a decision problem is a triple $\langle (\Omega, P), \mathcal{A}, U \rangle$ in which $\mathcal{A}$ is the set of actions, and U the utility function; Benz and van Rooij, 2007, 66). For example, take the very familiar example (17), with a toy model for a corresponding utility function.

---

[6]There may be some terminological confusion here with the notion of speech act conditionals / SACs that was mentioned above in relation to the work of Hesse et al. (2018). A SAC (maybe better called *biscuit conditional*) is a specific type of conditional sentence in which the consequent is factual, and does not depend on the truth of the *if*-clause. A conditional speech act / conditional assertion relates to the more general position that uttering a conditional sentence is not a single speech act that expresses a conditional proposition, but rather a combination of two speech acts. See Stalnaker (2011) for further discussion, cf. Ebert et al. (2014).

(17) Where can I buy an Italian newspaper?

P	$\frac{1}{3}$	$\frac{1}{3}$	$\frac{1}{3}$
W	w_1 : only @X	w_2 : only @Y	w_3 : @X&Y
a_X	6	0	4
a_Y	0	6	4

Here the set of actions $\mathcal{A} = \{a_X, a_Y\}$ (a_X: go to place X; a_Y: go to place Y) is fixed as part of the model. This is an unnatural assumption, because the asker of (17) most likely didn't have places X and Y in mind when she asked the question (otherwise something like 'Should I go to X or Y to buy an Italian newspaper?' would be more natural).

This problem also appears when a conditional answers a non-conditional question. For a decision problem representing a polar question $?q$, Alice will only distinguish between q-worlds and non-q-worlds for her utility function. She is not aware that the utility of her actions depends on p. In other words, Alice's U-function will not distinguish between a world w_1 in which p and $\neg q$ are true, and a world w_2 in which $\neg p$ and $\neg q$ are true. However, the conditional answer does distinguish between such worlds. In fact, the dependency between p and q is what is conveyed by the answer, and makes it relevant.

Switching decision problems Instead of the above-mentioned approach of applying existing utility-based frameworks to conditional propositions, I argue that conditional answers can be used to indicate that speaker A's original decision problem $?q$ can be reduced to a different decision problem $?p$ that is easier to resolve.

In (12), Alice's original decision problem was whether John came to party or not. Bob does not have full information to directly resolve A's decision problem, but does have conditional knowledge about a dependency between John's work and his coming to the party. By giving the conditional answer, he conveys to Alice that there is an alternative way to resolve her problem, namely by finding out about the progress of John's work.

The difference between (11) and (12) from section 2 can now be understood: whereas (12) allows to shift to an alternative decision problem that is (potentially) easier to resolve, finding out about the meteorological facts in B's answer in (11) is more difficult than the resolving the original problem of whether it rained. In other words, the answer in (11) invites A to shift to an alternative de-

cision problem that is more difficult than the original one, rendering it uncooperative.

Hence, the study of conditional answers argues for a dynamic turn in utility-based pragmatics of question-answer pairs: answers can lead to updating decision problems (and utility functions alongside) in the course of the conversation. As far as I know, such a dynamic take on utilities has not been proposed before. A way to formalize this idea is work in progress, by employing a probabilistic dynamic semantics (Yalcin, 2012), and combining probabilistic belief update with the utility function.

4 Conclusion

I have outlined various reasons for investigating conditional utterances from the perspective of answers to questions: understanding the use of conditionals in discourse, their information-structural properties, and the distribution of conditional perfection. Then I outlined some specific examples in which conditional answers are licensed, and argued that interlocutors have representations of each other's epistemic states. They update, and reason probabilistically about these representations in the course of the conversation. The utility of a conditional answer is measured by $P_B(K_A?p)$, but there is a second way in which the information conveyed by a conditional utterance counts as relevant, over and above the epistemic condition. Learning about the conditional dependency between p and q is relevant for A in the process of resolving her decision problem '$?q$', when this dependency allows her to switch from the original problem '$?q$' to an alternative decision problem '$?p$' that is easier to resolve.

References

Johan van der Auwera. 1997. Pragmatics in the last quarter century: The case of conditional perfection. *Journal of Pragmatics*, 27:261–274.

Johan van Benthem, Jelle Gerbrandy, and Barteld Kooi. 2009. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96.

Anton Benz and Robert van Rooij. 2007. Optimal assertions, and what they implicate. A uniform game theoretic approach. *Topoi*, 26:63–78.

Anton Benz and Jon Stevens. 2018. Game-theoretic approaches to pragmatics. *Annual Review of Linguistics*, 4:173–191.

Ingrid van Canegem-Ardijns. 2010. The indefeasibility of the inference that if not-A, then not-C. *Journal of Pragmatics*, 42(1):1–15.

Ingrid van Canegem-Ardijns and William van Belle. 2008. Conditionals and types of conditional perfection. *Journal of Pragmatics*, 40:349–376.

Fabrizio Cariani and Lance J. Rips. 2017. Experimenting with (conditional) perfection. Ms.

Ivano Ciardelli, Jeroen Groenendijk, and Floris Roelofsen. 2018. *Inquisitive semantics*. Oxford University Press.

Ivano A. Ciardelli and Floris Roelofsen. 2015. Inquisitive dynamic epistemic logic. *Synthese*, 192(6):1643–1687.

Benoît de Cornulier. 1983. *If* and the presumption of exhaustivity. *Journal of Pragmatics*, 7:247–249.

Christian Ebert, Cornelia Ebert, and Stefan Hinterwimmer. 2014. A unified analysis of conditionals as topics. *Linguistics & Philosophy*, 37(5):353–408.

Donka Farkas and Kim Bruce. 2010. On reacting to assertions and polar questions. *Journal of Semantics*, 27:81–118.

Marie-Christine Farr. 2011. Focus Influences the Presence of Conditional Perfection: Experimental Evidence. In Ingo Reich, Eva Horch, and Dennis Pauly, editors, *Proceedings of Sinn & Bedeuting 15*, pages 225–239. Universaar – Saarland University Press, Saarbrücken.

Gibson Ferguson. 2001. If you pop over there: a corpus-based study of conditionals in medical discourse. *English for Specific Purposes*, 20(1):61–82.

Kai von Fintel. 2001. Conditional strengthening: a case study in implicature. Ms., MIT.

Kai von Fintel. 2011. Conditionals. In Claudia Maienborn, Klaus von Heusinger, and Paul Portner, editors, *Semantics. An International Handbook of Natural Language Meaning. Volume 2*, pages 1515–1538. De Gruyter.

Michael Geis and Arnold Zwicky. 1971. On invited inferences. *Linguistic Inquiry*, 2(4):561–566.

Noah D. Goodman and Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, 20(11):818–829.

Elena Herburger. 2015. Conditional Perfection: the truth and the whole truth. In Sarah D'Antonio Mary Moroney and Carol Rose Little, editors, *Proceedings of SALT 25*, pages 615–635. LSA.

Christoph Hesse, Maryam Mohammadi, Maurice Langner, Judith Fischer, Anton Benz, and Ralf Klabunde. 2018. Communicating an understanding of intention: Speech act conditionals and modified numerals in a Q/A system. In Laurent Prévot, Magalie Ochs, and Benoît Favre, editors, *Proceedings of SemDIAL 2018 (AixDial)*, pages 103–111.

Michela Ippolito. 2013. Counterfactuals and Conditional Questions under Discussion. In Todd Snider, editor, *Proceedings of SALT 23*, pages 194–211. LSA.

James Isaacs and Kyle Rawlins. 2008. Conditional Questions. *Journal of Semantics*, 25:269–319.

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

Craige Roberts. 1996/2012. Information Structure: Towards an Integrated Formal Theory of Pragmatics. *Semantics & Pragmatics*, 5:6:1–69. Published version of a 1996 manuscript with the same title.

Robert van Rooij. 2004. Utility of Mention-Some Questions. *Research on Language and Computation*, 2:401–416.

Robert van Rooij and Katrin Schulz. 2019. Conditionals, causality and conditional probability. *Journal of Logic, Language and Information*, 28:55–71.

Robert Stalnaker. 2011. Conditional propositions and conditional assertions. In Andy Egan and Brian Weatherson, editors, *Epistemic Modality*, pages 227–248. Oxford University Press, Oxford.

Jon Scott Stevens, Anton Benz, Sebastian Reuße, and Ralf Klabunde. 2016. Pragmatic question answering: A game-theoretic approach. *Data & Knowledge Engineering*, 106:52–69.

Jos Tellings. 2016. *Counterfactuality in discourse*. Ph.D. thesis, UCLA.

Jos Tellings. 2020. Translation mining in the domain of conditionals: first results. Presentation at *Computational Linguistics in the Netherlands* (CLIN) 30.

Seth Yalcin. 2012. Context probabilism. In Maria Aloni, Vadim Kimmelman, Floris Roelofsen, Galit W. Sassoon, Katrin Schulz, and Matthijs Westera, editors, *Logic, language and meaning. 18th Amsterdam Colloquium*, pages 12–21. Springer.