

Probabilistic modeling of rational communication with conditionals *

Britta Grusdt
*Institut für Kognitionswissenschaft
Universität Osnabrück*

Daniel Lassiter
*School of Philosophy, Psychology and
Language Sciences
The University of Edinburgh*

Michael Franke
*Allg. Sprachwissenschaft &
Pragmatik
Eberhard Karls Universität Tübingen*

Submitted 2022-05-03 / First decision 2022-09-20 / Revision received 2022-09-28 /
Accepted 2022-10-03 / Published 2022-10-12 / Final typesetting 2023-11-17

Abstract While a large body of work has scrutinized the meaning of conditional sentences, considerably less attention has been paid to formal models of their pragmatic use and interpretation. Here, we take a probabilistic approach to pragmatic reasoning about indicative conditionals which flexibly integrates gradient beliefs about richly structured world states. We model listeners' update of their prior beliefs about the causal structure of the world and the joint probabilities of the consequent and antecedent based on assumptions about the speaker's utterance production protocol. We show that, when supplied with natural contextual assumptions, our model uniformly explains a number of inferences attested in the literature, including epistemic inferences, conditional perfection and the dependency between antecedent and consequent of a conditional. We argue that this approach also helps explain three puzzles introduced by Douven (2012) about updating with conditionals: depending on the utterance context, the listener's belief in the antecedent may increase, decrease or remain unchanged.

Keywords: indicative conditionals, Rational-Speech-Act model, computational pragmatics, causal Bayes nets, inferentialism

* We would like to thank three anonymous reviewers for their highly appreciated and very valuable feedback which helped to improve the manuscript substantially.

©2022 Britta Grusdt, Daniel Lassiter, Michael Franke
This is an open-access article distributed under the terms of a Creative Commons Attribution License (<https://creativecommons.org/licenses/by/3.0/>).

1 Introduction

Despite the long history of research on conditionals, no consensus has been reached on crucial questions concerning their semantics and pragmatics, let alone on fundamental questions such as whether conditionals have truth conditions at all. What is inarguably true is the diversity of how conditionals are interpreted, yet the circumstances that lead to different interpretations of conditionals depending, among others, on the utterance context are not well understood. An utterance of the conditional in (1) might be disillusioning when given by a professional epidemiologist, as one would naturally infer that the speaker is ignorant about the truth of the antecedent and the truth of the consequent. But how does this inference about the speaker’s uncertainty come about?

- (1) If the pandemic is over next year, the conference will take place offline again.

The conditional in (1) further suggests that the pandemic, mentioned in the antecedent, is related to the form of the conference, mentioned in the consequent. Does this inferred relationship originate from pragmatic reasoning or do we draw this inference because it is inherently part of the semantics of conditionals? The latter is advanced by prominent recent accounts of the meaning of conditionals, so-called *inferentialism* e.g., [Douven 2017](#), [Douven et al. 2018](#).

One long-standing problem that has particularly gained attention in relation with inferentialism is the problem of missing-link conditionals. Missing-link conditionals are — pragmatically infelicitous — conditionals with no discernible causal or conditional relationship between antecedent and consequent, as in example (2).

- (2) If Jo doesn’t pass the exam, it will rain tomorrow.

Contrary to inferentialist accounts, we argue that it is possible to explain the observation that missing-link conditionals are infelicitous solely on the basis of pragmatics, making use of a notion of *listener surprise* about the speaker’s utterance choice.

In this paper, we explore a formalization of interlocutors’ pragmatic reasoning about the use and interpretation of conditionals within the Rational-Speech-Act (RSA) modeling framework ([Frank & Goodman 2012](#), [Franke &](#)

Jäger 2016, Goodman & Frank 2016).¹ We aim to contribute to the understanding of how the inferences commonly associated with conditionals arise and how different utterance contexts influence the interpretation of conditionals. In a nutshell, a probabilistic pragmatics approach enables us to combine two aspects, which we here argue are key to obtaining a more realistic picture of the pragmatics of conditionals: pragmatic reasoning in the vein of Grice (2013) and richly structured, probabilistic world knowledge, particularly about *a priori* plausible causal relationships between antecedent and consequent.

To illustrate this, consider example (1) again. We argue that the inferences outlined above can be explained by pragmatic reasoning about the speaker's utterance choice: a cooperative speaker who is in the position to assert the antecedent and/or the consequent alone, should make this stronger claim in order to be as informative as required (Quine 1965, Grice 2013, 1989). Given that the speaker did not claim the truth of the antecedent nor the consequent, a listener would plausibly infer uncertainty on the side of the speaker about whether the antecedent and the consequent are true. It is in this sense that we expect Gricean pragmatic reasoning about the speaker's protocol of choosing utterances (e.g., governed by *Maxims of Conversation* or similar) to matter for the interpretation of conditionals.

As to the effects of prior world knowledge, the conditional in (1) also illustrates that the listener's prior beliefs may influence the strength of the beliefs the listener holds after receiving information in form of a conditional. When an assertion of (1) is accepted, the listener's degree of belief in the conference taking place offline conditional on the pandemic being over will be high, simply as a result of the conditional's semantic meaning. But the degree to which the listener believes in the conference *not* taking place offline conditional on the pandemic *not* being over (arguably one possible measure of a Conditional Perfection reading, see below) crucially depends on the listener's further beliefs about the world. The important, though perhaps uncontroversial, point to notice in connection with this example is that prior world knowledge is likely to impact the beliefs the listener holds after accepting and interpreting an utterance of a conditional sentence.

At large, it is challenging to account for how people update beliefs in light of new conditional information whereas there is a widely adopted general ac-

¹ When we write 'conditional', we refer to indicative conditionals whose antecedent and consequent are simple propositions, to which we restrict the considered kind of conditionals in this paper.

count of how people update beliefs in light of new factual information, *probabilistic conditioning*. In this approach, factual information is represented by a probability distribution over a set of possible worlds. When a new piece of information E is received, an updated belief state is formed by assuming that the information is true and adjusting the distribution accordingly. The consequence is that $P(A)$, the prior belief distribution on proposition A , is updated to a posterior distribution $P(A | E)$ so that:

$$P(A | E) = \frac{P(A \wedge E)}{P(E)} .$$

This idea lies at the heart of Bayesian models in epistemology and decision theory (e.g., Earman 1992) as well as computational cognitive science (Tenenbaum et al. 2011). Conditional sentences present a challenge to the generality of the Bayesian approach to learning, because there is no consensus about how the update should work for conditionals, or whether conditioning is even applicable e.g., Grove & Halpern 1997, Douven & Romeijn 2011. Many authors have argued that conditionals cannot be regarded as having truth conditions at all (for a recent review see Rothschild 2015), but in its standard form, Bayesian update cannot be applied to sentences without truth-conditions. Even if conditionals do have truth-conditions, it remains unclear how to implement a plausible update operation for conditional information.

In an influential paper, Douven 2012 argues that there is no uniform way to capture the update effect of conditionals within Bayesian theory, regardless of the assumed semantics. The argument hinges on the relationship between the listener’s prior and posterior beliefs about the antecedent. Douven shows that — depending on the context in which a conditional is uttered — the inferred probability of the antecedent may either decrease, increase or remain unchanged compared to the listener’s beliefs prior to the uptake of the conditional. Surveying a variety of possible semantics for conditionals, he shows that all of them predict that Bayesian update should have a uniform effect on the probability of the antecedent. We argue that it is possible to explain the non-uniform update effect of conditionals, that we will refer to as *Douven’s puzzle*, by standard Bayesian belief revision if we also take into account pragmatic reasoning and world knowledge. Specifically, prior knowledge about how the relevant propositions are likely to relate to each other influence assumptions about what a rational speaker would say under various circumstances. Assumptions about a speaker’s production protocol in turn influence the listener’s interpretation of an utterance of a conditional.

From this perspective, Douven’s examples make a strong case for the importance of pragmatic and contextual information in the interpretation of conditionals, and provide a useful test suite for our model which attempts to formalize these factors.²

We argue that prior knowledge about the likely causal relationship between antecedent and consequent is not only important to account for Douven’s puzzle, but generally mediates a listener’s interpretation of a conditional. For instance, in (1) our background knowledge supports the intuitive inference that there is a causal and probabilistic relationship between the form of a conference and the state of the pandemic. However, no such relationship seems reasonable in missing-link conditionals like (2). But why does background knowledge affect our interpretation in this way, leading to a sense of infelicity when a dependency between consequent and antecedent is implausible? We argue that assumptions about the causal relationship between antecedent and consequent have an (implicit) influence on the speaker’s utterance choice and thereby also on the listener’s interpretation, assuming a Gricean pragmatic interpretation of conditionals uttered by a cooperative speaker. As we will show, Gricean pragmatic reasoning — formalized in the Rational-Speech-Act model and supplemented by structured representations of causal world knowledge — is better able to account for missing-link conditionals than inferentialists have claimed. Overall, the key to the success of our model stems from leveraging the natural stochastic association between (i) whether antecedent and consequent are causally related and (ii) whether “If A , C ” (henceforth $A \rightarrow C$) is the most informative true utterance the speaker can make.

The paper is structured as follows. Section 2 introduces the vanilla RSA model and details our implementation of the model for reasoning about the use of conditionals. Section 3 and 4 describe how this model of pragmatic reasoning helps explain key phenomena of interest, on the basis of unspecific as well as contextualized, concrete examples, including Douven’s puzzles, conditional perfection readings, and general inferences about the causal relationship between antecedent and consequent. Section 5 discusses how the presented approach positions itself with respect to missing-link and biscuit conditionals. Section 6 reflects critically and concludes.

² Like-minded proposals to account for Douven’s puzzle have been advanced recently e.g., Günther 2018, Eva, Hartmann & Rad 2019, Vandenburg 2020, making use of the different intuitive causal structures of Douven’s scenarios. However, none of these proposals considers the crucial role of pragmatic reasoning.

2 A Rational-Speech-Act model for communication with conditionals

To provide a principled and uniform formal treatment of pragmatic reasoning which explains the observations raised in the introduction, we turn to probabilistic pragmatic modeling. The Rational-Speech-Act (RSA) model is a prominent instance of a formalization of Gricean pragmatic reasoning using tools from probability calculus, decision and game theory (see [Franke & Jäger 2016](#), [Goodman & Frank 2016](#): for an overview). RSA models are probabilistic, data-driven and based on the assumption that linguistic behavior is goal-oriented and in this regard (approximately) optimal. A noteworthy benefit of a probabilistic and data-driven approach is that the predictions by the models can often be directly compared to quantitative aspects of experimental data, thereby allowing the statistical comparison of theoretically relevant models based on empirical data (e.g. [Qing & Franke 2015](#), [Degen et al. 2020](#), [Franke & Bergen 2020](#)). Moreover, as probabilistic modeling is prominent in other areas in the cognitive sciences, it becomes relatively easy to integrate insights from these other areas as well, such as belief formation ([Goodman & Stuhlmüller 2013](#), [Herbstritt & Franke 2019](#)), sequential adaptation ([Schuster & Degen 2019](#)), or learning biases ([Brochhagen, Franke & van Rooij 2018](#)). Since it was first introduced by [Frank & Goodman \(2012\)](#) in the context of a language reference game, the number of RSA models and modeled phenomena has been growing; scalar implicatures ([Goodman & Stuhlmüller 2013](#)), hyperboles ([Kao et al. 2014](#)) as well as more complex phenomena such as the interpretation of vague or polite language ([Lassiter & Goodman 2017](#), [Yoon et al. 2016](#)), projective content ([Qing, Goodman & Lassiter 2016](#)), metaphors ([Kao, Bergen & Goodman 2014](#)) or social meaning ([Burnett 2019](#)) have been investigated by means of RSA.³

There are several reasons why RSA also seems promising for modeling the interpretation of conditionals. First of all, we expect an influence of the availability of non-conditional utterances that the speaker might have chosen instead of the conditional. In this way, RSA formalizes a Gricean account of listeners' interpretations and the speakers' utterance choices. Since RSA is able to flexibly integrate contextual knowledge, we can model quite different utterance contexts, including situations in which richly structured world

³ A hands-on introduction to RSA modeling, using the probabilistic programming language WebPPL ([Goodman & Stuhlmüller 2014](#)), that we also use to implement our model, is provided by [Scontras, Tessler & Franke \(2017\)](#).

knowledge is relevant. This will be crucial in our account of the effects of causal knowledge on the interpretation of conditionals.

In the following, Section 2.1 gives a general introduction to the vanilla version of the RSA model. Subsequent sections then elaborate on the specific adaptations necessary in order to capture reasoning about the use of conditionals against the background of richly structured causal world knowledge. We use Section 2.3 to demonstrate the main underlying ideas of the model by means of a toy example before we lay out the assumed priors on the generalized set of states in the sections thereafter.

2.1 The vanilla Rational-Speech-Act model

At the heart of the vanilla RSA model lies the formalization of a cooperative Gricean speaker who, when trying to communicate a state s , probabilistically selects an utterance u by preferably choosing utterances that are not only true, but also maximize the amount of relevant information conveyed to a literal listener. The pragmatic listener is modeled as a rational interpreter who combines prior beliefs with the speaker’s protocol of choosing utterances by using Bayes’ rule.

In order to capture the pragmatic speaker’s behavior, in particular, in order to ground out a notion of truth and informativity, the RSA model considers a literal listener first, whose interpretation behavior is defined as a conditional probability distribution $P_{\text{lit}}(s \mid u)$ obtained by updating any prior beliefs $P_{\text{prior}}(s)$ about likely world states s with the set of states $\llbracket u \rrbracket$ where utterance u is permissible.⁴

$$[1] \quad P_{\text{lit}}(s \mid u) \propto \delta_{s \in \llbracket u \rrbracket} \cdot P_{\text{prior}}(s)$$

The pragmatic speaker is then defined in terms of a notion of *utterance utility* $U(u; s)$, reflecting how informative u is for communicating s .⁵

$$[2] \quad U(u; s) = \log P_{\text{lit}}(s \mid u)$$

⁴ The formula in Equation [1], and in Equations [3] and [4] for other model components, gives the probability up to proportionality (\propto), leaving the normalizing constant of the probability distribution implicit. If $F(x) \geq 0$ is a non-normalized score for any $x \in X$ with X a finite set, the notation $P(x) \propto F(x)$ is shorthand for $P(x) = \frac{F(x)}{\sum_{x'} F(x')}$. Moreover, $\delta_{s \in \llbracket u \rrbracket}$ is the Kronecker delta function, which returns 1 when its Boolean argument — the denotation function of utterance u applied to state s — is true, otherwise 0.

⁵ A common modification of the basic definition of utterance utility in terms of informativity further accounts for differences among utterances by using utterance cost regarding, for instance, their complexity, saliency or social compatibility (e.g., see Qing, Goodman & Lasnik 2016, Gates et al. 2018). We do not include utterance costs in our model, but they could easily be integrated.

The informativity of an utterance u as description of a state s is defined as the log-likelihood (negative surprisal) of the literal listener’s beliefs for state s after hearing utterance u . The probability that a speaker in state s will choose utterance u is then defined as a soft-max operation on utility scores:

$$[3] \quad P_S(u | s) \propto \exp(\alpha \cdot U(u; s))$$

α is a free model parameter governing how closely the speaker approximates utility maximization: the higher α , the more likely the speaker is to choose the utility-maximizing utterance. At the extremes, a hyperrational speaker ($\alpha \rightarrow \infty$) would only choose utterances that maximize utility, while a randomizing speaker ($\alpha = 0$) would choose randomly among true utterances.

Finally, the pragmatic listener’s interpretation is captured by a conditional probability distribution, $P_{PL}(s | u)$, which represents the listener’s *a posteriori* beliefs (after having heard utterance u) about the probability of state s , taking the priors over states and a Gricean speaker’s utterance-choice behavior into account.

$$[4] \quad P_{PL}(s | u) \propto P_S(u | s) \cdot P_{\text{prior}}(s)$$

2.2 World states, utterances & assertability

World states. Conditionals like $A \rightarrow C$ are often associated with the speaker’s uncertainty about whether A and/or C are true. Therefore, to model pragmatic reasoning about conditionals, we include potentially uncertain speakers into our modeling. Concretely, we look at the partition of possible worlds into the four types of worlds which agree on the truth values of A and C : $W = \{w_\emptyset, w_A, w_C, w_{AC}\}$. The concrete set of states used are probability distributions over worlds $w \in W$.

There are at least two prominent possibilities of how to interpret probability distributions: as precise *objective* chances or as *subjective* beliefs. Consequently, world states in our model can be interpreted in different ways, too. For one, we can think of $s \in S$ as the true beliefs of a maximally competent speaker about the objective chance of each type of world. In this case, the conversational goal, implicitly defined in the vanilla RSA model, is to communicate the true (objective, but intrinsically stochastic) world state known to the speaker. For another, we can conceptualize world states as a representation of an uncertain speaker’s subjective beliefs $s \in S$ about the true (non-probabilistic) state of the world $w \in W$. Under this interpretation, the

vanilla RSA model implicitly treats the conversational goal as that of communicating the speaker’s belief state (see Aloni 2007, Franke 2011). The interpretation of world states, as either objective or subjective, may matter for the interpretation of the assertability conditions presented below. If not stated otherwise, we refer to the subjective interpretation of model states.

Utterance alternatives. Predictions of Gricean pragmatic reasoning strongly depend on the assumed set of alternative utterances. There has been much discussion of alternative sets for scalar items like *some*, *warm* and *or* (e.g. Matsumoto 1995, Katzir 2007), but much less for pragmatic reasoning about conditionals (for some discussion see van der Auwera 1997, von Stechow 2001). The selection of alternative utterances considered in our model is largely governed by the desire to present a balanced set of alternative utterances sufficient to describe the most salient differences in the set of world states. Utterances are compositionally built from literals, possibly negated. They may be combined with *and* to form a conjunction, with *if* to form a conditional or with the word *likely*. Table 1 lays out all alternative utterances together with the rule used to compute the update effects of each — its “assertability condition”, to be introduced next.

utterance type	assertability in state s	example:	
		utterance u	assertability u in state s
conjunction	$P^{(s)}(\phi, \psi) \geq \theta$	$A \wedge \neg C$	$P^{(s)}(A = a, C = \neg c) \geq \theta$
literal	$P^{(s)}(\phi) \geq \theta$	A	$P^{(s)}(A = a) \geq \theta$
conditional	$P^{(s)}(\psi \phi) \geq \theta$	$A \rightarrow \neg C$	$P^{(s)}(C = \neg c A = a) \geq \theta$
likely + literal	$P^{(s)}(\phi) > 0.5$	<i>likely</i> $\neg C$	$P^{(s)}(C = \neg c) > 0.5$

Table 1 Types of utterances with corresponding assertability conditions and an example, ordered from most informative utterance on top to least informative at the bottom. For conditionals and conjunctions, $\phi \neq \psi$.

Assertability conditions. Given a world state $s \in S$ and an utterance $u \in U$, the definition of the semantics $\llbracket u \rrbracket \subseteq S$ serves as the anchoring of pragmatic reasoning in literal interpretation, defined in Equation [1]. As especially the semantic meaning of conditionals is highly controversial, we would like

to stay as uncommitted and encompassing as possible. This is possible, to a certain extent, if we focus not on the nature of the denotation function $\llbracket u \rrbracket$ but rather at the functional role it plays in the architecture of the pragmatic reasoning model. In particular, since the utility function in Equation [2] and the speaker rule in Equation [3] entail that whenever $s \notin \llbracket u \rrbracket$, the speaker will not choose u when in state s , the main effect of $\llbracket u \rrbracket$ is to give *assertability conditions* and — as a side-effect — information about how informative each utterance is.⁶ Consequently, our model lays out a general method of computing update effects of conditionals at the level of a literal interpreter with the ulterior goal of defining reasonable speaker behavior, while avoiding as much as possible concrete commitment to a specific semantic interpretation.

We treat utterances of literals like “ A ” as conveying that the target state s makes the probability that A is true (i.e., $A = a$) high enough for conversational purposes; the corresponding probability, e.g., $P^{(s)}(A = a)$ for literal “ A ”, must exceed a certain threshold θ for the respective utterance to be assertable.⁷ In a model of objective chance, as we assume on an objective interpretation of the probability distributions building up the set of world states, determinate truth corresponds to probability 1. That is, a speaker will treat A as true just in case A is determinately true in s — when $P^{(s)}(A = a) = 1$. The assumption that speakers sometimes assert things that are not certain, but very likely true, yet justifies a threshold below 1 as assertability condition. We also make use of this assumption on a subjective interpretation of world states; a factual sentence A is thus assertable as long as the speaker’s subjective belief in A is sufficiently large ($P^{(s)}(A = a) \geq \theta$, respectively $P^{(s)}(A = \neg a) \geq \theta$ when $u = \neg A$).

Similarly, “*likely A*” directly conveys that the subjective probability of A in s is greater than 0.5 as we assume that “*likely A*” is assertable in s iff $P^{(s)}(A = a) > 0.5$. This aspect of our account is reminiscent of expressivist

⁶ The informativity of the modeled utterances generally follows the order shown in Table 1. As we were pointed at by an anonymous reviewer, depending on the chosen prior over states, it is, however, possible for a conditional (e.g., $A \rightarrow C$) to be *literally* more informative than a literal (e.g., C) since the assertability of a literal does not per se entail the assertability of a conditional, e.g. $P^{(s)}(c) \geq \theta \not\Rightarrow P^{(s)}(c | a) \geq \theta$. This is an interesting observation, in particular in the context of conditionals whose consequent is independent of the antecedent as for instance in concessive conditionals (e.g., “*Even if ...*”) which is certainly worth looking at in future work, but beyond the scope of what we can cover in this paper.

⁷ We write $P^{(s)}$ to refer to probabilities *within* states to distinguish them from probabilities *across* states; $P^{(s)}(X)$ denotes the probability assigned to any event X that may be inferred from state s , e.g., $P^{(s)}(A = a)$ is the probability of A to be true in s .

accounts of probability language (Yalcin 2012, Moss 2015, Swanson 2015). On the objective interpretation, this means that *likely* expresses high objective chance, consistent with the empirical findings of Ülkümen, Fox & Malle (2016) and Lassiter (2018). Note, however, that these authors show that *likely* can also express subjective uncertainty.

In parallel fashion, we render the assertability conditions of an indicative conditional $A \rightarrow C$ as $P^{(s)}(C = c \mid A = a) \geq \theta$, corresponding either to a high objective chance of C to be true when A is true or to a strong belief in this conditional probability (on part of the speaker), assuming an objective and subjective interpretation respectively. Since we do not want or need to commit to a specific semantic theory of conditionals here, we settle for motivating this condition as a plausible minimal bound on assertability that should be acceptable from a wide range of theoretical perspectives. For theories that are able to support the equation $P(A \rightarrow C) = P(C \mid A)$ while avoiding triviality results (e.g., Lewis 1976, Hájek 1989), the derivation is strictly parallel to the factual case above. This includes non-propositional theories (Edgington 1995, Bennett 2003), trivalent theories (de Finetti 1936, Milne 1997, Lassiter 2020) and various others (van Fraassen 1976, Stalnaker & Jeffrey 1994, Kaufmann 2004, Khoo 2016, among others). Our assertability condition is also particularly natural for theories that render the truth-condition of $A \rightarrow C$ as $P(C \mid A) = 1$ (e.g., Moss 2015).

The status of our assertability condition for conditionals is somewhat murkier from the perspective of other prominent theories such as Stalnaker (1968) and Kratzer (1991), as well as strict conditional theories. Because of the complexity of the way that they assign truth-values to epistemically possible worlds where the antecedent A is false, these theories can make $P(A \rightarrow C)$ high even while $P(C \mid A)$ is low. As a result, our assertability condition is stronger than these accounts would predict. However, we note that there is by now an enormous body of empirical evidence supporting the equation between the probability of a conditional and the corresponding conditional probability (Hadjichristidis et al. 2001, Evans & Over 2004, Douven & Verbrugge 2010, among many others). This evidence problematizes a key prediction made by the latter group of theories: that a conditional can be judged highly probable simply because of the likely falsehood of its antecedent. Instead, situations where the antecedent is false are generally judged irrelevant to the probability of a conditional, in a probabilistic analogue of the classic paradoxes of the material conditional (Edgington 1995). We do not doubt that the theories under consideration have theoretical resources available that

may allow them to avoid this problem—for example, by using pragmatic reasoning to explain why false-antecedent cases are not considered relevant in assertion (Lewis 1976, Grice 1989). But doing so would be tantamount to adopting our assertability condition or something quite close to it. Therefore, we believe that our results should be relevant to theorists with a wide variety of semantic commitments, including those for whom probabilistic reasoning has not previously played a major theoretical role.

2.3 Toy example

Let us consider a toy example to illustrate model predictions, and to motivate further generalizations to be introduced hereafter. Concretely, we will consider a case with just three world states, all equally likely and just constructed for the sake of illustration. The topic of conversation in this example is whether Alex and Chris are likely to go to a party. The states we look at differ in the probability they assign to all four logical possibilities of Alex and/or Chris going to the party.

- (s_1) Alex and Chris both really like to go out. Both are seen at most parties, but whether either comes is unrelated to whether the other comes.
- (s_2) Alex and Chris go slightly more often than not, but usually not without each other (e.g., they are a couple, best friends, etc.).
- (s_3) Alex and Chris each go out more often than not, but have no connection with each other.

These three scenarios are translated into the probability distributions $P(A, C)$ shown at the top of Table 2 where A (C) denotes whether Alex (Chris) comes to the party.

With this context in mind, imagine the following utterance Bigi directs at Wobo:

- (3) Bigi: *Chris will come to the party.* [“ C .”]

Given that Bigi only says what she believes to be true and s_1 , s_2 and s_3 describe all possible scenarios, Wobo should infer that Bigi describes the situation modeled in s_1 since this utterance (C) is neither assertable in s_2 nor s_3 .

Now, imagine Bigi to utter the following conditional instead:

- (4) Bigi: If Alex comes to the party, Chris comes too. [“If A , C .”]

As noted in Table 2, $A \rightarrow C$ is assertable in states s_1 and s_2 . Therefore, under a literal interpretation of Bigi’s utterance, s_1 and s_2 are considered equally likely ($P_{\text{lit}}(s_1 | u = A \rightarrow C) = P_{\text{lit}}(s_2 | u = A \rightarrow C) = 0.5$). However, under a pragmatic interpretation of $A \rightarrow C$, s_2 is judged as more likely than s_1 ($11/16$ vs. $5/16$). This is because Bigi could have chosen a more informative utterance to communicate s_1 .

	s_1	c	$\neg c$	s_2	c	$\neg c$	s_3	c	$\neg c$
(a) $P(A, C)$	a	0.81	0.09	a	0.6	0.05	a	0.36	0.24
	$\neg a$	0.09	0.01	$\neg a$	0.05	0.3	$\neg a$	0.24	0.16
	$u = \text{likely } C$	1			1			1	
(b) $u = A \rightarrow C$		1			1			0	
$u = C$		1			0			0	
$u = A \wedge C$		0			0			0	
(c) $P_{\text{lit}}(s_i u = \text{likely } C)$		1/3			1/3			1/3	
$P_{\text{lit}}(s_i u = A \rightarrow C)$		1/2			1/2			0	
$P_{\text{lit}}(s_i u = C)$		1			0			0	
$\sum_{u'} P_{\text{lit}}(s_i u')$		11/6			5/6			1/3	
(d) $P_S(u = \text{likely } C s_i)$		$\frac{1/3}{1/6} = 2/11$			$\frac{1/3}{5/6} = 2/5$			$\frac{1/3}{1/3} = 1$	
$P_S(u = A \rightarrow C s_i)$		$\frac{1/2}{1/6} = 3/11$			$\frac{1/2}{5/6} = 3/5$			0	
$P_S(u = C s_i)$		$1/11 = 6/11$			0			0	
(e) $P_{\text{PL}}(s_i u = \text{likely } C)$		$\frac{2}{11}/(\frac{2}{11} + \frac{2}{5} + 1) = 10/87$			$\frac{2}{5}/(\frac{2}{11} + \frac{2}{5} + 1) = 22/87$			$1/(\frac{2}{11} + \frac{2}{5} + 1) = 55/87$	
$P_{\text{PL}}(s_i u = A \rightarrow C)$		$\frac{3}{11}/(\frac{3}{11} + \frac{3}{5}) = 5/16$			$\frac{3}{5}/(\frac{3}{11} + \frac{3}{5}) = 11/16$			0	
$P_{\text{PL}}(s_i u = C)$		$\frac{6}{11}/\frac{6}{11} = 1$			0			0	

Table 2 Model predictions for the scenario from Section 2.3 where $\alpha = 1, \theta = 0.9$. (a) States s_1, s_2, s_3 . (b) Assertability of utterances given states with respectively most informative utterances in bold. (c) Literal interpretation. (d) Speaker production likelihoods. (e) Pragmatic interpretation.

2.4 Inferring latent causal relations

The example in Table 2 shows how the model introduced so far yields the inference that $A \rightarrow C$ is most likely associated with s_2 , by rather straightforward Gricean reasoning. It also demonstrates how, on top of inferring a state s , Wobo might draw inferences about the likely causal relation between A and C which may have led to Bigi’s beliefs as captured in s_1, s_2

or s_3 . By motivation of the example, the coming of Alex and Chris was assumed to be *independent* in states s_1 and s_3 , while the very reason for writing a table like in s_2 was because we assumed that there was a stochastic relationship between Alex’s and Chris’ coming to the party. Suppose that there are two equally likely relations $r \in \{\text{independent}, \text{dependent}\}$, meaning that A and C are either independent or dependent. Intuitively, $P(S = s_2 \mid r = \text{independent})$ is much smaller than $P(S = s_1 \mid r = \text{independent})$ or $P(S = s_3 \mid r = \text{independent})$, and also $P(S = s_2 \mid r = \text{dependent})$ is much higher than $P(S = s_1 \mid r = \text{dependent})$ or $P(S = s_3 \mid r = \text{dependent})$. In this way, by Bayesian inference, we can obtain an indirect inference of a likely causal/stochastic relation between A and C just from probabilistic pragmatic reasoning *and* natural assumptions about the differential likelihood between different causal/stochastic relations r and states s . Note the importance of the pragmatic reasoning for drawing an inference about the likely relation in this example: under a literal interpretation of the conditional $A \rightarrow C$, Wobo would not show a preference between s_1 , where A and C are likely dependent, and s_2 , where they are likely independent.

We can think of this model as a sequence of inferences: beliefs about r stochastically inform s via world knowledge or intuitions about dependence/independence and s , in turn, stochastically informs the speaker’s utterance choice u due to pragmatic constraints on what counts as a good utterance. Schematically: $P(r) \Rightarrow P(s \mid r) \Rightarrow P(u \mid s)$ which is depicted in Figure 1. We can then derive $P(r \mid u)$ via Bayes’ rule. In the following we specify and motivate a concrete prior structure, in particular for the $P(s \mid r)$ part, so as to be able to derive general predictions from this model for what we may consider a *default context*, where no specific world knowledge is assumed to be available regarding antecedent and consequent.




Figure 1 Through world knowledge, beliefs about the causal relation $r \in R$ inform probabilities $s \in S$ which are, in turn, considered as known by the speaker. Probabilities in s then directly influence the speaker’s utterance choice $u \in \mathcal{U}$.

2.5 Prior over world states in default context

A state s where variables A and C are assumed to be independent ($r = A \perp\!\!\!\perp C$) is represented by probabilistically independent distributions where $P^{(s)}(A, C) = P^{(s)}(A) \cdot P^{(s)}(C)$. Since in the default context, we do not make any specific assumptions, A and C are both assigned a uniform prior probability over the interval $[0, 1]$:

$$[5] \quad P^{(s)}(A = a), P^{(s)}(C = c) \sim \text{Uniform}(0, 1)$$

Together with the assumption of independence, this yields the following probability distribution over partitions of possible worlds, representing a single state s when $r = A \perp\!\!\!\perp C$:

$$\begin{aligned} P(w_{AC}) &= P^{(s)}(A = a) \cdot P^{(s)}(C = c), & P(w_A) &= P^{(s)}(A = a) - P(w_{AC}) \\ P(w_C) &= P^{(s)}(C = c) - P(w_{AC}), & P(w_{\emptyset}) &= 1 - (P(w_{AC}) + P(w_A) + P(w_C)) \end{aligned}$$

To derive the probability distributions over partitions of possible worlds that represent states where variables A and C are assumed to be dependent, we distinguish between two possible *types* of causal relation: either A has causal power to provoke C ($R = A \rightsquigarrow C$) or vice versa ($R = C \rightsquigarrow A$).⁸ While the causal relation, R , merely provides information about the causal direction, the concrete *instances* of a causal relation, denoted as r , are distinguished based on how exactly they affect the conditional probabilities, as shown in Table 3. $R = A \rightsquigarrow C$, for instance, tells us that the outcome of variable A has direct influence on the outcome of variable C , and the instance of the causal relation, r , further tells us *how* the outcome of A influences the outcome of C . For example, we write $r = A \rightsquigarrow^{++} C$ to pick out the class of probability distributions where the truth of A increases the probability that C is true, whereas $r = A \rightsquigarrow^{--} C$ picks out probability distributions where the *falsity* of A increases the probability that C is true. Putting this differently, when $r = A \rightsquigarrow^{++} C$, w_{AC} will be considerably more likely than w_A ($P(w_{AC}) \gg w_A$), and when $r = A \rightsquigarrow^{--} C$, $P(w_C) \gg P(w_{\emptyset})$.

⁸ The choice to model only a single relevant cause is corroborated by many empirical studies that have shown people's tendency to neglect alternative causes (e.g., see Fernbach & Darlow 2010, Fernbach, Darlow & Sloman 2011, Fernbach & Rehder 2013, Haggmayer & Waldmann 2007, Krynski & Tenenbaum 2007); also, in a theoretical model from Icard & Goodman (2015), the loss of information resulting from neglecting alternative causes was on average too small as to justify their consideration.

causal relation (R)	instance causal relation (r)	interpretation
$A \rightsquigarrow C$	$A \overset{++}{\rightsquigarrow} C$	Truth of A increases probability for truth of C
	$A \overset{-+}{\rightsquigarrow} C$	Falsity of A increases probability for truth of C
$C \rightsquigarrow A$	$C \overset{++}{\rightsquigarrow} A$	Truth of C increases probability for truth of A
	$C \overset{-+}{\rightsquigarrow} A$	Falsity of C increases probability for truth of A

Table 3 Notation for dependent causal relations (types and instances). The instance of the causal relation provides information about the associated joint probability tables, spelled out in column ‘interpretation’.

Formally, the probability distributions of the dependent states can be described as leaky-noisy or model with binary variables allowing for positive as well as negative causes. Noisy-or models (Pearl 1988) describe the relationship between an effect variable and its cause variables where each cause is capable of producing the effect independently of all other causes; corresponding to a logical OR-function where the effect is only true if at least one of its causes is true. Leaky noisy-or models (Díez 1993) comprise “background noise” represented by an additional cause variable that is always present and summarizes all potential causes of the effect that are not explicitly modeled. Figure 2 shows a graphical representation of the leaky noisy-or model underlying the dependent states in our model with the corresponding conditional probabilities of the effect C to be true when A is true or false respectively, assuming non-deterministic relations where variables A and B have causal power τ , respectively β , to provoke C . Since we want to stick to the simplest set of induced probability distributions, but still need to cover all possible stochastic relations between the two explicitly modeled variables A and C to have a balanced set of states that preserves the informativity of utterances, we use a generalization of the classical noisy-or model which further allows negative causes: not only the truth, but also the falsity, of a cause can have an influence on the truth of the effect (e.g., see Hyttinen, Eberhardt & Hoyer 2011).⁹ Therefore, C is likely true when A is false and a negative cause

⁹ If we only used positive causes, states where $P(w_{AC})$ and $P(w_{\emptyset})$ tend to be low would be very rare and thus, conditionals like $A \rightarrow \neg C$ would only be assertable in very few states

($r = A \overset{-+}{\rightsquigarrow} C$) or when A is true and a positive cause ($r = A \overset{++}{\rightsquigarrow} C$); in both cases the truth of C may independently be due to background noise B .




Figure 2 Graphical representation of a leaky noisy-or model (left) with a single explicitly modeled cause A and the corresponding conditional probabilities of C when A is true or false (right). Variable B summarizes all potential other causes of C ; τ and β denote the causal power of A , respectively B , to induce the truth of C .

To instantiate the set of probability distributions for the respective dependent causal relations, we specify the prior distributions over the respective causal power (τ), noise (β) and prior probability of the parent variable (v_p) as shown in Figure 3.¹⁰ The values of the hyperparameters for the beta distributions of the causal power and noise are chosen such that the mean of the distribution of the causal power τ exceeds the assertability threshold, set to 0.9 in all simulations below; for the causal power of the background noise, the parameters are simply reversed such that the prior distribution $P(\beta)$ is skewed towards 0. As we consider a default context here, where no further information is available, the prior probability of the parent variable (v_p) is sampled from a uniform distribution over the interval $[0, 1]$. The joint probability distributions over A, C are then derived based on τ, β and v_p . Depending on the relation R and the cause variable being a positive or a negative cause, β, v_p and v_c correspond to different (conditional) probabilities, as listed in Table 4 where v_c denotes the conditional probability of the effect to be true when the cause variable is true: $v_c = \tau + \beta - \tau \cdot \beta$ (see Figure 2).

which would, in turn, render this utterance very informative. As a result, this particular conditional would become more informative than for instance literals, which does not seem reasonable. Using states with positive and negative causes results in a balanced set of states such that any conjunction is more informative than any literal and any literal is, in turn, more informative than any conditional.

¹⁰ Note that even though the notion of *causal* relations is eventually not important as the modeling hinges on certain probabilistic dependencies between events (depending on the relation), it is important for the choice of our prior distributions which is motivated by relations that are causal by nature.

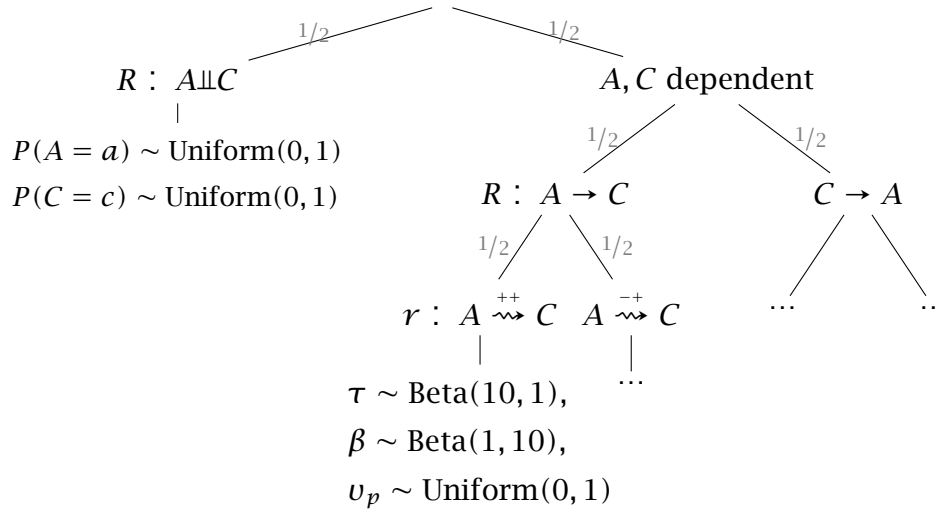


Figure 3 Graphical representation of the procedure for sampling a state s from the prior in the default context.

instance causal rela- tion (r)	u_p	u_c	β
$A \overset{++}{\rightsquigarrow} C$	$P^{(s)}(A = a)$	$P^{(s)}(C = c \mid A = a)$	$P^{(s)}(C = c \mid A = \neg a)$
$A \overset{-+}{\rightsquigarrow} C$	$P^{(s)}(A = \neg a)$	$P^{(s)}(C = c \mid A = \neg a)$	$P^{(s)}(C = c \mid A = a)$
$C \overset{++}{\rightsquigarrow} A$	$P^{(s)}(C = c)$	$P^{(s)}(A = a \mid C = c)$	$P^{(s)}(A = a \mid C = \neg c)$
$C \overset{-+}{\rightsquigarrow} A$	$P^{(s)}(C = \neg c)$	$P^{(s)}(A = a \mid C = \neg c)$	$P^{(s)}(A = a \mid C = c)$

Table 4 Probabilities (u_p, u_c, β) that define the joint probability distribution of a state s , $P^{(s)}(A, C)$, for each instance of a dependent causal relation. u_p is the prior probability of the cause, u_c is the conditional probability of the effect to be true when the cause is true and β is the power of the unmodeled variables to provoke the effect corresponding to the conditional probability of the effect to be true when the explicitly modeled cause is false.

The joint probability distributions over partitions of possible worlds are then computed as follows when $R = A \rightsquigarrow C$ (similarly for $R = C \rightsquigarrow A$):

$$\begin{aligned} P(w_{AC}) &= P(c \mid a) \cdot P(a), & P(w_A) &= P(\neg c \mid a) \cdot P(a) \\ P(w_C) &= P(c \mid \neg a) \cdot P(\neg a), & P(w_\emptyset) &= P(\neg c \mid \neg a) \cdot P(\neg a) \end{aligned}$$

In total, we sample 10,000 probability distributions that comprise the set of world states used by our model. More concretely, we first sample a causal relation r from its prior distribution given in Equation [6]. Our choice to put a prior on the causal relation, which determines the shape of the associated sampled probability distributions, is primarily based on work from [Tenenbaum & Griffiths \(2003\)](#) and [Griffiths & Tenenbaum \(2005\)](#); they were the first who put priors on the causal structure itself to predict human causal judgments instead of focusing on learning the causal strength of a, possibly non-existent, causal link (e.g., see [Cheng 1997](#)).

$$[6] \quad P(r) = \begin{cases} 1/2 & \text{if } r = A \parallel C \\ 1/8 & \text{if } r \in \{A \rightsquigarrow C, A \rightsquigarrow\!\!\!\!\rightsquigarrow C\} \\ 1/8 & \text{if } r \in \{C \rightsquigarrow A, C \rightsquigarrow\!\!\!\!\rightsquigarrow A\} \end{cases}$$

Based on the causal relation r , we then sample probability distributions according to the procedure described above. A visualization of the sampled states is given in Figure 4 which shows histograms of the probabilities for each of the four possible worlds across all sampled states for three selected causal relations.¹¹

2.6 Communicating causal information implicitly via conditionals

So far, we have specified how the joint probability distributions are derived under the assumption of particular causal relations. That is, strictly speaking, our world states have two components, a probability distribution s and a causal relation r . Independently of the causal relation from which s originates, the choice probabilities of a speaker who aims to communicate her beliefs about the world can then be written as:

$$[7] \quad P_S(u \mid r, s) = P_S(u \mid s) \propto \exp(\alpha \cdot (\log P_{\text{lit}}(s \mid u)))$$

where $P_{\text{lit}}(s \mid u) = \sum_{r'} P_{\text{lit}}(r', s \mid u)$

¹¹ The source code and all modeling results are publicly available: https://osf.io/6bshq/?view_only=1703a3417a1343a5a66b78ac8ce206c2.

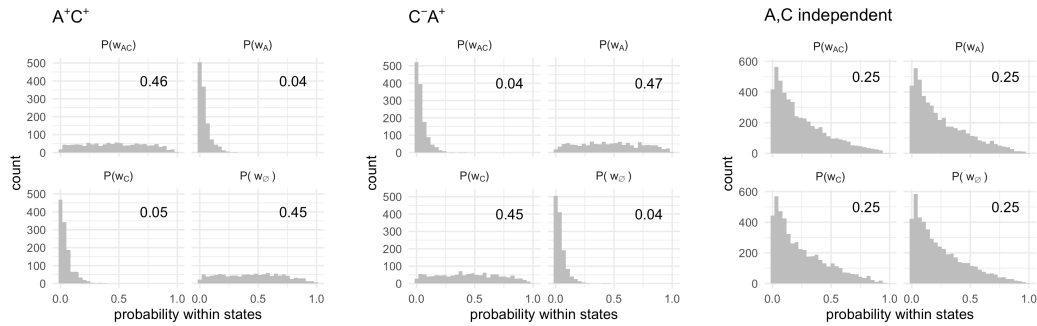


Figure 4 Histograms of the probabilities of the four possible worlds, w_{AC} , w_A , w_C and w_{\emptyset} of all sampled probability tables $P^{(s)}(A, C)$ with $r = A \rightsquigarrow^+ C$ (left), $r = C \rightsquigarrow^- A$ (middle) and $r = A \perp\!\!\!\perp C$ (right). Numbers in the upper right corners are the expected values for the respective worlds.

The speaker’s goal of communication, when using conditional sentences, that we assume in this paper, is first and foremost to convey their beliefs about the antecedent and the consequent. In other words, we start by exploring a probabilistic model of communication with conditionals from the most austere assumption, namely that not only the assertability of a conditional does *not* hinge on any putative causal relation necessarily, but that also the purpose of communication itself is *not* to directly communicate information about the causal relation.¹²

Before we discuss the results of our simulations, we would like to add a final note on the interpretation of the world states in our model. Instead of considering them as simple pairs consisting of a probability distribution and a causal relation, we can also think of them as *causal Bayesian networks* (Pearl 1988, 2009, 1988), henceforth abbreviated as ‘Bayes nets’. They have, similar to the closely related formalism of Structural Causal Models, a rich tradition of supporting semantic theories of conditionals already, (e.g. Hiddleston 2005, Pearl 2009, 2013, Rips 2010, Briggs 2012, Kaufmann 2013, Lucas & Kemp 2015, Santorio 2016, Lassiter 2017). Bayes nets represent sets of variables (here binary variables) and the dependencies among them; they consist of a directed, acyclic graph that defines the relations among variables, and a set of conditional probabilities for each variable given all possible instantiations

¹² There are circumstances where the goal of the communication reasonably includes the causal relation which we leave for future work here.

of all its *direct* parent nodes, which simplify to unconditional probabilities for variables without parent nodes. This set of (conditional) probabilities is sufficient to define the joint probability distribution over all variables represented in the graph; see Figure 5 for a Bayes net that represents the joint probability distribution s_2 from Table 2.

In sum, the model we explore here can be seen as capturing the implicit communication of causal information by (i) treating world states as causal Bayes nets, (ii) identifying the purpose of utterance (the *question under discussion* (Roberts 1996) or the relevance projection (Kao et al. 2014)) to be the precise communication of the probability table s (the speaker’s beliefs about joint truth of A and C) and (iii) a pragmatic process of utterance generation favoring true and informative utterances.




Figure 5 Bayes net representing s_2 from Table 2 consisting of a graphical representation (left) and a set of associated (conditional) probabilities (right) that define $P^{(s_2)}(A, C)$.

3 Pragmatic inferences from conditionals in default contexts

In this section, we explore model predictions in default contexts. We can think of these as the model’s predictions generalized over a wide range of more specific contexts, or, relatedly, as the model’s predictions for the interpretation of utterances of conditionals in unbiased, out-of-the-blue contexts. We will particularly look at the listener’s inferences from an utterance of a conditional about the speaker’s uncertainty about A and C , the speaker’s beliefs about any systematic relation between A and C , and the strength of a conditional perfection reading. Doing so, we show how informativity-driven pragmatic choice of utterances leads to *de facto* assertability conditions as

postulated by inferentialist accounts without having to stipulate these directly.¹³

Setting the scene. Let us reconsider example (4), repeated here from above, now uttered in an out-of-the-blue context.

- (4) Bigi: If Alex comes to the party, Chris comes too. ["If A, C."
Wobo: Who are these guys? What party are you talking about?

Even without any strong prior convictions, Wobo is likely to draw pragmatic inferences from Bigi's utterance. Intuitively, Wobo would infer that Bigi is uncertain about whether Alex comes to the party, similarly for Chris, and that Chris' coming to the party is not entirely unrelated to Alex's coming. Moreover, the implicit dependency between antecedent and consequent may be interpreted to be so strong that the conditional is understood as biconditional; it does not seem unnatural for Wobo to infer that, according to Bigi, Chris comes to the party *only if* Alex comes.

To analyse the listener's *a posteriori* beliefs and the speaker's utterance choices we make use of the following definitions concerning the modeled states. The set of states in which the speaker is *uncertain about event X* contains all and only states s such that the probability of X in s , $P^{(s)}(X)$, is neither too low nor too high:

$$[8] \quad \text{Uncertain}(X) = \{s \mid 1 - \theta \leq P^{(s)}(X) \leq \theta\}$$

A similar construction captures the speaker's *certainty about whether X is true*:

$$[9] \quad \text{Certain}(X) = \{s \mid P^{(s)}(X) > \theta\} \cup \{s \mid P^{(s)}(X) < 1 - \theta\}$$

The uncertainty / certainty about X given Y is captured analogously.

Hyperrational utterance choices in default contexts. According to the model presented here, the listener's inferences about the speaker's epistemic state can be put down, at least in part, to a Q-implicature (Atlas & Levinson 1981, Horn 1984): from the fact that the speaker decided *not* to utter a more specific and thus more informative utterance than $A \rightarrow C$, the pragmatic listener should deem it unlikely that the speaker refers to a state s in which a more informative utterance would also apply. To see whether such an alternatives-based explanation is endorsed by the present modeling

¹³ All results reported below were obtained with the literal meaning threshold $\theta = 0.9$, and the rationality parameter $\alpha = 3$. Qualitatively identical results were obtained for parameters from a grid with $\theta \in [0.9, 0.95, 0.975]$, $\alpha \in [1, 3, 5, 10]$.

setup, we first look at a speaker who always chooses the utterance with the highest utility (hyperrational speaker where rationality parameter $\alpha \rightarrow \infty$). Figure 6 shows model predictions from the speaker’s point of view for different probabilistic beliefs with respect to propositions A and C given the causal relation between antecedent and consequent. It clearly shows that a (hyperrational) speaker chooses utterances in dependence of her belief state: when the speaker is uncertain about both propositions, the best utterance is either a conditional or “*likely* Φ ” which are the two least informative utterance types.¹⁴ On the other hand, the speaker’s best utterance is either a conjunction or a literal when she is certain about both propositions, that is, when she is certain that A is true (or false) and certain that C is true (or false), but she is not necessarily certain that a conjunctive event occurs (e.g. $A = a \wedge C = \neg c$). This is also the reason, why conjunctions do not seem to be preferred over literals when A and C are independent and the speaker is certain about A and certain about C (Figure 6 (i)). In this case, when the speaker’s best utterance is a literal, there simply is no conjunction that truthfully describes the given state (e.g., $P^{(s)}(a, c) = 0.82, P^{(s)}(a, \neg c) = 0.08, P^{(s)}(\neg a, c) = 0.08, P^{(s)}(\neg a, \neg c) = 0.02$). When the speaker is only certain about one proposition, her best utterance is a literal (Figure 6 (iii)).

Figure 6 (ii) further reveals that, when the speaker is uncertain about both A and C , her utterance choice is strongly influenced by their causal relation: for these states the speaker’s best utterance will always be “*likely* Φ ” when A and C are independent, whereas when A and C are dependent, it will almost certainly be a conditional.

Inferences about causal dependency. Since pragmatic interpretation is here modeled as backwards-inference based on the speaker’s utterance choice protocol, we can already anticipate from the above results of hyper-rational speakers that pragmatic interpreters may draw rather specific inferences about the causal relation between A and C from an utterance of $A \rightarrow C$. In what follows, we look at inferences about the causal relationship in more detail, assuming “normal speakers” ($\alpha = 3$).

Figure 7 shows the causal inferences that listeners draw about A and C when the speaker utters the conditional $A \rightarrow C$. We observe that *a posteriori* the listener (literal and pragmatic) assigns very low probability to states where $r = A \overset{+}{\rightsquigarrow} C$ or $r = C \overset{+}{\rightsquigarrow} A$, as these are very unlikely to give rise to a probability table in which the conditional $A \rightarrow C$ is assertable.

¹⁴ For simplicity, we discuss all results assuming an epistemic interpretation of the probabilities in our world states, but an interpretation based on objective chance is equally applicable.

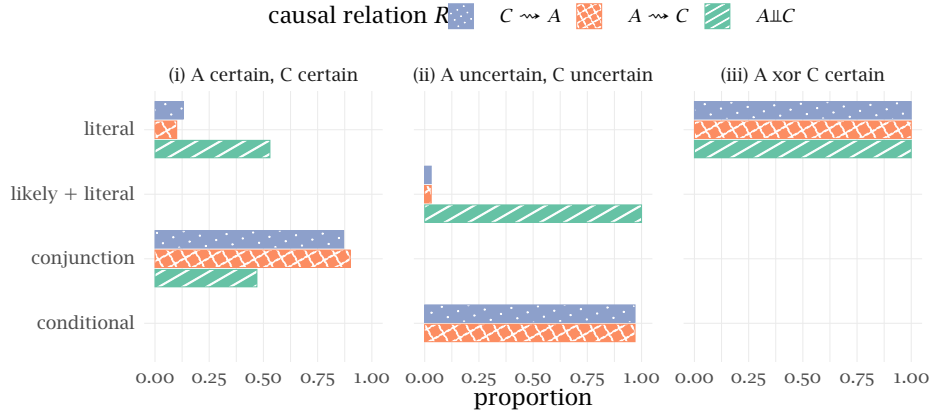


Figure 6 Relative frequency of how often each utterance type is the speaker’s *best* choice for a set S of 10,000 states sampled from the prior (default context), given that the speaker is (i) certain or (ii) uncertain about A and about C , i.e., $\forall s \in S : s \in \text{Certain}(A) \wedge s \in \text{Certain}(C)$, respectively $\forall s \in S : s \in \text{Uncertain}(A) \wedge s \in \text{Uncertain}(C)$, or (iii) the speaker is uncertain about the truth of one proposition but certain about the truth of the other, e.g., $\forall s \in S : s \in \text{Uncertain}(A) \wedge s \in \text{Certain}(C)$.

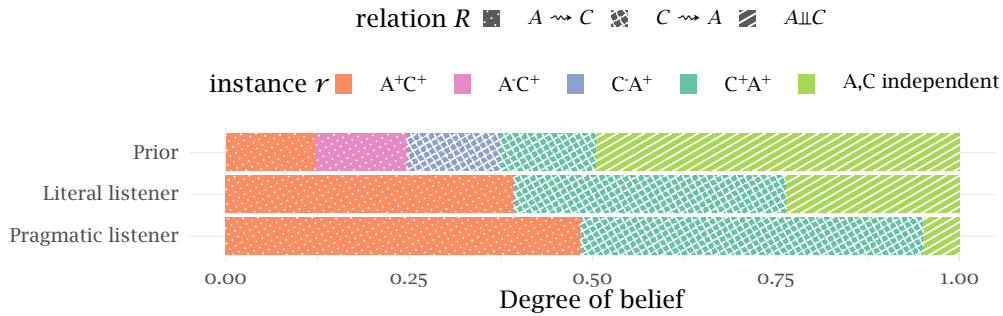


Figure 7 Degree of belief in each causal relation (A^+C^+ is shorthand for $A \overset{++}{\rightsquigarrow} C$, analogous for other relations) at all three levels of interpretation; prior to uptake of the conditional $A \rightarrow C$ and *a posteriori* given a literal / pragmatic listener.

Interestingly, the listener is not committed to a single underlying causal relation, but instead merely infers that there is a positive dependency between antecedent and consequent: A tends to be true (false) when C is true (false) and vice versa. The pragmatic listener assigns almost the entire probability mass to the corresponding causal relations ($A \rightsquigarrow^{++} C, C \rightsquigarrow^{++} A$), the literal listener approximately 75%.

This result suggests that also under a pragmatic interpretation, the listener needs further knowledge to disambiguate the underlying causal structure since, in this most general context, it is not possible to infer that the antecedent is a cause of the consequent or that, as under a diagnostic reading of the conditional $A \rightarrow C$, the consequent is a cause of the antecedent.

Another interesting result concerns the states where A and C are independent: even though the literal listener largely diminishes her beliefs in the independence of A and C as compared to her beliefs prior to the speaker's utterance of the conditional $A \rightarrow C$, the pragmatic listener assigns a still smaller probability to states where antecedent and consequent are causally independent. This is an important and interesting result that bears emphasis. Since the pragmatic listener combines a rich representation of the speaker's beliefs about truth of propositions and their causal relation with Gricean pragmatic reasoning, the pragmatic listener concludes more about the (speaker's beliefs about the) causal structure of the world than is entailed by the semantics of a conditional. This additional causal-pragmatic inference essentially rides piggyback on standard Gricean Quantity reasoning.

To see this, let us consider the perspective of the hyperrational speaker again (Figure 6). Those states where A and C are independent are exclusively states for which an informative Gricean speaker will either prefer to utter a bare proposition (conjunction or literal) which is more informative than a conditional or "*likely* Φ ". As a result, the pragmatic listener infers a causal relation from entirely unbiased assertability conditions for conditionals and standard Gricean Quantity reasoning. Causal inference comes up as a pragmatic inference without having to stipulate an additional pragmatic constraint concerning a (causal) relation between propositions, let alone hard-coding such a requirement in the semantic meaning.

The strength of conditional perfection readings. In the example given in the beginning of this section, it does not seem astounding to interpret the "if" in "*If Alex comes to the party, Chris comes too*" as "*if and only if*". It might even seem to be a quite acceptable, if not natural, inference although we did

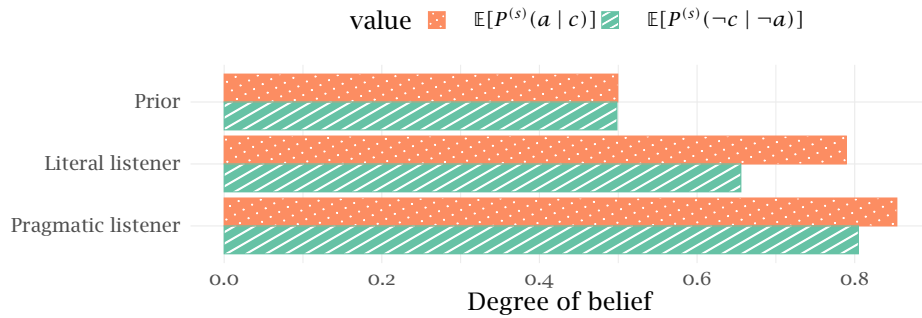


Figure 8 Degree of belief in the two CP-related conditional probabilities $P^{(s)}(-c | \neg a)$ and $P^{(s)}(a | c)$; prior to the uptake of the conditional $A \rightarrow C$ and *a posteriori*, given a literal or pragmatic listener.

not specify any further context and although, from a logical point of view, this inference is *not* valid.

The phenomenon that conditionals are sometimes interpreted as biconditionals, known as *conditional perfection* (CP), has caught much attention in the literature, especially since Geis & Zwicky 1971. CP remains a topic of ongoing debate; no consensus has for instance been found concerning its prevalence or the circumstances that trigger a CP reading (e.g. see Newstead 1997, von Stechow 2001, Oberauer & Wilhelm 2003, Moldovan 2013).

Neither exists a standard measure that quantifies the degree to which a conditional receives a CP reading. Here we will refer to two inferences that are prominently considered in the literature on conditional reasoning, ‘*Denying the antecedent*’ (DA) and ‘*Affirming the consequent*’ (AC), shown below. The endorsement of DA or AC inferences suggests that participants interpret the conditional as biconditional since only then these are logically valid inferences (e.g., Newstead 1997, Evans et al. 2007).

DA: $A \rightarrow C. \neg A. \therefore \neg C.$

AC: $A \rightarrow C. C. \therefore A.$

To learn how strongly $A \rightarrow C$ is interpreted as biconditional, we will therefore look at the listener’s expected beliefs (about the speaker’s beliefs) $P^{(s)}(-c | \neg a)$ and $P^{(s)}(a | c)$. Only when “*if*” is interpreted as “*iff*” the two considered quantities should be large.

As can be seen in Figure 8, without any further contextual assumptions our model predicts a quite strong CP-reading for the interpretation of con-

ditionals in the default context. The speaker's utterance of $A \rightarrow C$ elicits an increase in the listener's beliefs (about the speaker's beliefs) in the conditional probabilities $P^{(s)}(\neg c \mid \neg a)$ and $P^{(s)}(a \mid c)$ as compared to her prior beliefs. This is true for the literal and the pragmatic interpretation, yet the pattern is more pronounced in the latter. This result is to a great extent due to the representation of dependent world states as noisy-or models (see Section 2.5). In this way, this interpretation is explained as something akin to an I-implicature (Atlas & Levinson 1981, Levinson 2000), as suggested by Horn (2000); CP-readings are supported in large part by what is arguably a cognitively economic, perhaps stereotypical representation format of causal dependency between events. However, we also see that pragmatic reasoning about alternatives further strengthens the CP-reading quantitatively, similar to the accounts of van der Auwera (1997) and von Stechow (2001) which share with our model the feature that CP is a result of a listener reasoning about the speaker's production protocol — given von Stechow's (2001) account, including, crucially, the possibility that the speaker could simply have asserted the consequent.

Deriving inferentialist assertability conditions. So far, we showed that the current modeling setup predicts that speakers will use conditional sentences predominantly in cases where there is a (causal/inferential) relationship between antecedent and consequent and that listeners, therefore, infer such a relationship from an utterance of a conditional.¹⁵ These predictions particularly challenge the idea advanced by advocates of inferentialism that a (causal/inferential) relation between antecedent and consequent is part, in whatever form, of the core semantics of conditionals. We here argue that an austere assertability condition for conditionals in combination with rich representations of contextual (causal) world knowledge and pragmatic reasoning is sufficient to derive the kind of assertability conditions postulated by inferentialist accounts. In particular, we will consider the assertability condition formulated in Equation [10] which was proposed by van Rooij & Schulz

¹⁵ Our model predicts the dependency relation to be a defeasible inference, but it is not predicted to arise in any context (e.g., see the discussion on missing-link and biscuit conditionals in Section 5). Therefore, it is also compatible with Lassiter's (2022) account of when the dependency between antecedent and consequent will (or will not) arise based on discourse coherence.

(2019) improving on like-minded work of Douven 2008.

[10] $A \rightarrow C$ is acceptable/assertable only if

$$\Delta^*P = \frac{P(c | a) - P(c | \neg a)}{1 - P(c | \neg a)} \geq \theta.$$

The general idea behind assertability criteria of this kind is that an utterance of a conditional $A \rightarrow C$ is acceptable only if $P(c | a)$ is high (like we assume as well here) and, in addition, $P(c | \neg a)$ is low, that is, C being true is only likely when A is true.

We do not argue that Equation [10] necessarily holds, empirical research is needed to find out whether the acceptance/assertability of conditionals is related to such criteria.¹⁶ Instead, we aim to investigate how the model that we propose here relates to accounts that propose this kind of assertability conditions for conditionals, possibly bringing Inferentialists' ideas and pragmatic reasoning closer together.

We find that our pragmatic model for the use of conditionals derives the criterion from Equation [10], in the sense that whenever the model predicts a (hyperrational) speaker to use a conditional in some state s , the value Δ^*P calculated from the probability table entailed by s is indeed very high. To see this, Figure 9 shows the distribution of Δ^*P measures for three sets of states: (i) sampled from the prior (default contexts), (ii) sampled from the prior conditioned on $A \rightarrow C$ being assertable (literal speaker condition), and (iii) the subset of states sampled in the literal speaker condition, in which a hyper-rational speaker would utter $A \rightarrow C$ (pragmatic speaker condition). Figure 9 reveals, reassuringly, that Δ^*P is not always high for any state sampled from the prior. It also shows that just from our austere assertability condition alone, namely $P^{(s)}(c | a) \geq \theta$, the average associated Δ^*P increases, even if there are still quite a number of cases where what we may call a "literal speaker" might use $A \rightarrow C$ while the measure Δ^*P is quite low. This clearly shows that, despite biases for "simple situations" introduced by *noisy-or* parameterization of state priors, the assertability condition that $P^{(s)}(c | a) \geq \theta$ alone does not guarantee that Δ^*P is high. But for the pragmatic speaker, we see that the Δ^*P measure is exclusively very high, thus lending support to

¹⁶ While some of the empirical studies that have been conducted to date suggest that the relationship between antecedent and consequent has an influence on the assertability of conditionals (e.g., see Douven & Verbrugge 2012, Skovgaard-Olsen, Singmann & Klauer 2016), others found no such effect (e.g., see Singmann, Klauer & Over 2014, Oberauer, Weidenfeld & Fischer 2007).

A probabilistic model of communication with conditionals

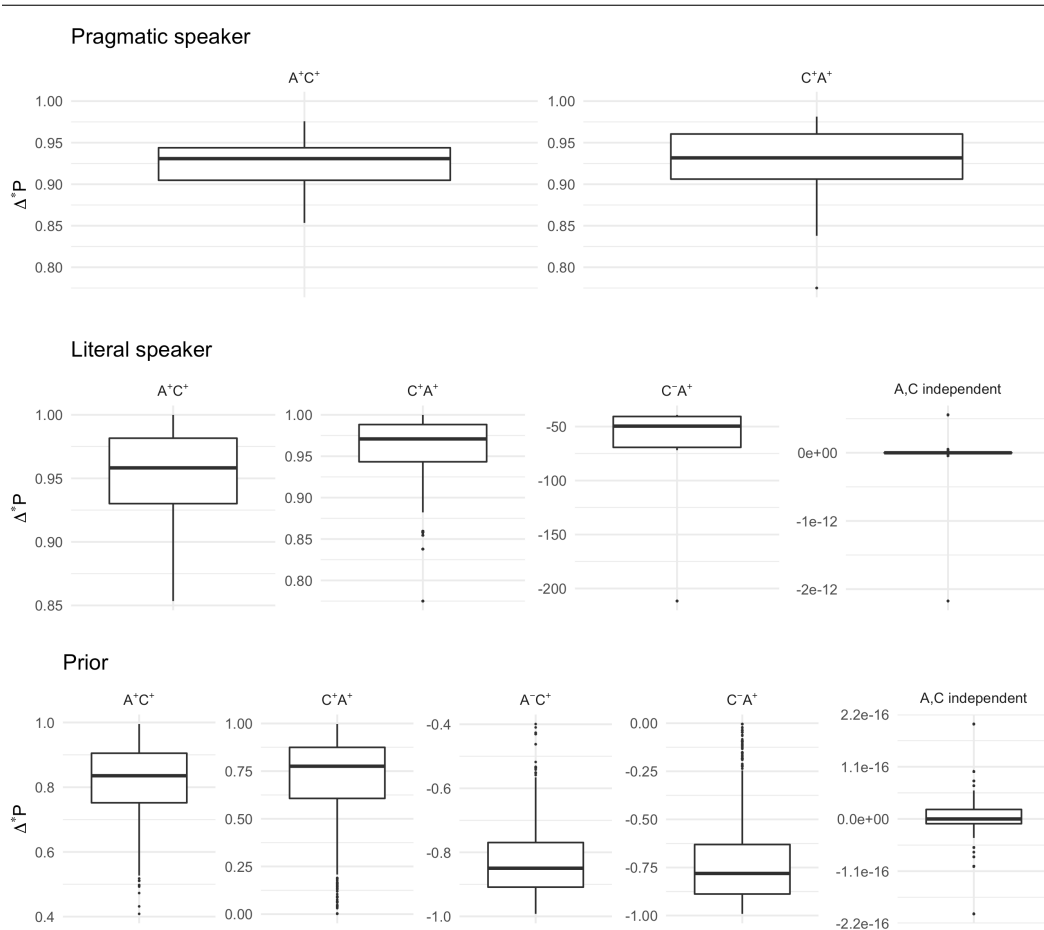


Figure 9 Distribution of Δ^*P values for 10,000 states randomly sampled from (i) the prior (bottom), (ii) from the prior given $A \rightarrow C$ is assertable (middle), and (iii) for states from (ii) where additionally $A \rightarrow C$ is the pragmatic speaker's best choice (top). Causal relations are abbreviated, e.g., A^+C^+ is shorthand for $A \rightsquigarrow^{++} C$.




Figure 10 $\Delta^*P^{(s)}$ values, zoomed into the range where $0 \leq \Delta^*P^{(s)} \leq 1$ for s from a set of 10,000 states sampled from the prior given $A \rightarrow C$ is assertable (literal speaker condition in Figure 9) where $A \rightarrow C$ is *not* the best choice of a hyperrational speaker ($\alpha = \infty$).

the idea that a straightforward model of Gricean pragmatic use of conditionals explains an otherwise stipulative assertability condition for conditionals.

We saw that when $A \rightarrow C$ is asserted by our pragmatic speaker, Δ^*P is large, but is it also the case that a large value of Δ^*P is sufficient for a hyperrational pragmatic speaker to utter the conditional $A \rightarrow C$? If it was a sufficient condition, our pragmatic speaker should prefer this conditional whenever Δ^*P is large. Figure 10 shows boxplots in the range $0 \leq \Delta^*P \leq 1$ of the Δ^*P -values of states where $A \rightarrow C$ is assertable, but *not* the most likely utterance for a hyperrational speaker.¹⁷ The utterance type of the hyperrational speaker’s best choice for the respective states — a literal, a conjunction or a conditional other than $A \rightarrow C$ — is shown on the x-axis; color codes represent the causal relation of the states. Clearly, there are states where $\Delta^*P^{(s)}$ is large and the speaker chooses a different utterance than $A \rightarrow C$. Particularly interesting for us are those states where the speaker does not choose a different conditional but an utterance that is more informative than a conditional, that is, a conjunction or a literal. These are situations where the predictions of our model diverge from the predictions of accounts arguing that a conditional is a assertable/acceptable when Δ^*P is large: although

¹⁷ For states where $R = C \rightsquigarrow A$, Δ^*P has a minimum value of -212 when the hyperrational speaker’s best utterance is a literal, and a minimum value of -69.3 when it is a conjunction. When $R = A \perp C$, $\Delta^*P^{(s)}$ clusters closely around 0.

$A \rightarrow C$ is (literally) assertable, the speaker may have good reasons to choose a different utterance (to be maximally informative) and so, the conditional $A \rightarrow C$ might still be rejected. This explains why a criterion like the one defined in Equation [10] might be a necessary, but not a sufficient condition for the assertability of a conditional.

Overall the predictions from our model corroborate the hypothesis that the dependency relation does not need to be incorporated into the semantics of conditionals. The model shows that pragmatic reasoning about the speaker's conditional utterance $A \rightarrow C$ is a possible way to derive the dependency relation—given an appropriate representation of the modeled variables, in particular their underlying causal structure.

4 Interpretation in concrete contexts: Douven's puzzle

So far we showed that an RSA model combined with unbiased default priors is able to explain general pragmatic inferences associated with utterances of conditionals. In the following we will turn towards predictions for utterances of conditionals in rather specific and explicitly given contexts of use. In particular, we investigate how the present setup helps explain a puzzle put forward by Douven (2012). Douven contrasts three cases of conditionals uttered in concrete contexts, that we will fully cite below. In these contexts, learning a conditional $A \rightarrow C$ either leads to an increase (the Skiing Example given in (5), discussed in Section 4.1) or a decrease (the Garden Party Example given in (6), discussed in Section 4.2) in the listener's degree of belief in the truth of A or the listener's degree of belief in the antecedent A does not change at all (the Sundowners Example given in (7), see Section 4.3).

We argue here that an explicit representation of contextually-grounded world knowledge that does not only comprise plausible prior beliefs about the probability of the represented variables, but also caters for their causal structure, is sufficient to explain how pragmatic listeners adjust their beliefs about the antecedent in one way or another after receiving information in form of a conditional. The key to the explanation we propose here is that, in each case, the listener learns a piece of causal information, that is, the listener learns that propositions are causally related where this was previously not expected or deemed rather unlikely.¹⁸

¹⁸ A perhaps more realistic picture would be to model a listener as completely *unaware* of the causal relation in question; that agent's stock of explicitly entertained alternatives would

4.1 The skiing case

The Skiing Example is a case where, intuitively, the listener’s degree of belief in the antecedent increases.

The Skiing Example (Douven 2012). Harry sees his friend Sue buying a skiing outfit. This surprises him a bit, because he did not know of any plans of hers to go on a skiing trip. He knows that she recently had an exam and thinks it unlikely that she passed. Then he meets Tom, another friend of Sue’s, who is just on his way to Sue to hear whether she passed the exam, and who tells him:

- (5) If Sue passed the exam, her father will take her on a skiing vacation.
 \rightsquigarrow listener belief in antecedent *increases*

In this example, there are three relevant propositions: E (Sue passed the exam), S (Sue goes skiing) and C (Sue buys skiing clothes). The listener Harry has observed C , so his beliefs in C are high, possibly 1. The speaker Tom utters the conditional $E \rightarrow S$. We want to explain how this can lead to an increase in Harry’s probabilistic beliefs about E .

Our explanation hinges on assuming that Harry has certain plausible beliefs about the propositions involved and their causal relationship. It proceeds in three steps:

- i. From pragmatic reasoning, the listener infers from the utterance of $E \rightarrow S$ that the speaker likely believes in a causal relation $E \rightsquigarrow^{++} S$: Sue passing the exam increases her chance to go on a skiing trip.
- ii. The listener takes the speaker to be an authority on the matter and, at least to a certain extent, also increases degrees of beliefs in the causal relation $E \rightsquigarrow^{++} S$.
- iii. Since the listener also has a high degree of belief in C , and given that it is plausible to assume that in general a relation of the kind $S \rightsquigarrow^{++} C$ holds, the listener ends up with a higher degree of belief in E after processing the utterance.

just not represent that contingency. Since adding agent’s unawareness to a model of pragmatic reasoning is possible but technically rather involved (e.g. Heifetz, Meier & Schipper 2006, Franke & de Jager 2011, Franke 2014), this paper makes the simplifying assumption that the listener is aware of the possible causal connection but deems it very unlikely to begin with.

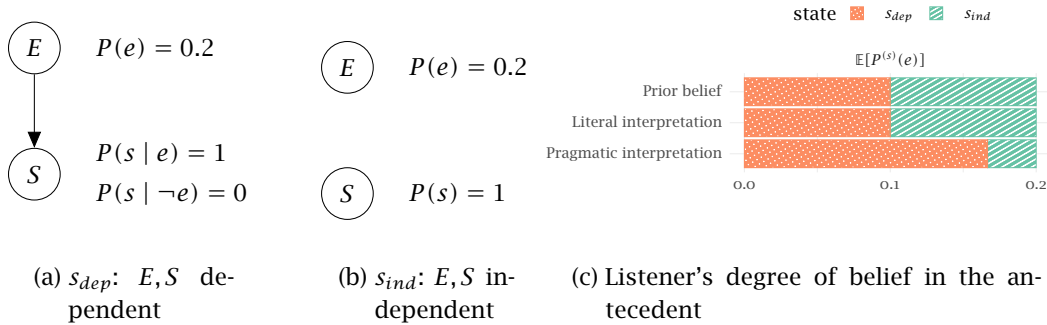


Figure 11 Bayes nets and results for the pragmatic reasoning part in the Skiing Example with $\mathcal{U} = \{S, \text{likely } S, E \rightarrow S\}$, $\alpha = 1$, E : pass exam, S : go skiing. Both states, s_{dep} and s_{ind} , are assigned equal prior probability.

To illustrate this reasoning schema, we offer one concrete context model for the Skiing Example in Figure 11.¹⁹ There are two Bayes nets that are in line with the speaker's utterance $E \rightarrow S$, one in which passing the exam stands in a direct causal relation to going on a skiing trip (Figure 11(a)), and one in which it does not (Figure 11(b)).²⁰

Their entailed joint probability distributions are spelled out in Table 5. The situation model in Figure 11 gives the listener's (Harry's) beliefs about what Tom might believe about the relation between E and S . Given the context-story, there is no indication that Harry believes that Tom believes that the dependent Bayes net is any more or less likely than the independent one. Therefore, we assign equal prior probability to both Bayes nets, even though this is not crucial for the case at hand.

¹⁹ A more realistic choice than using a single independent state with $P^{(s_{ind})}(s) > \theta$, would be to include several independent states, e.g. with $P^{(s)}(s) \in [0, 0.5, 1]$ which would, however, require a larger set of utterances such that for every state there is at least one utterance assertable. For the sake of a simpler discussion, here and in the following examples, we give just one set of concrete numbers of the relevant (conditional) probabilities.

²⁰ Even though it seems less probable, there is the possibility that the causal relation between E and S is reversed (i.e., $S \rightarrow E$ instead of $E \rightarrow S$); Sue may for instance study extra hard *because* her father invites her to go on a skiing trip. Note that in this case, the listener's observation of Sue buying skiing clothes would, however, not increase the listener's degree of belief in Sue passing the exam.

(a) $P^{(s_{dep})}(E, S)$

	s	$\neg s$
e	$P(e) \cdot P(s e) = 0.2$	$P(e) \cdot P(\neg s e) = 0$
$\neg e$	$P(\neg e) \cdot P(s \neg e) = 0$	$P(\neg e) \cdot P(\neg s \neg e) = 0.8$

(b) $P^{(s_{ind})}(E, S)$

	s	$\neg s$
e	$P(e) \cdot P(s) = 0.182$	$P(e) \cdot P(\neg s) = 0.018$
$\neg e$	$P(\neg e) \cdot P(s) = 0.728$	$P(\neg e) \cdot P(\neg s) = 0.072$

Table 5 Entailed joint probability distributions of variables E, S of the two Bayes net s_{ind}, s_{dep} shown in Figure 11; E : ‘pass exam’, S : ‘go on skiing trip’.

With the set of alternative utterances $\mathcal{U} = \{S, \text{likely } S, E \rightarrow S\}$, $\alpha = 1$ and the probabilities and causal relations as shown in Figure 11, we get the following probability distribution for the pragmatic listener when $u = E \rightarrow S$: $P_{PL}(s = s_{dep} | u = E \rightarrow S) = 5/6$ and $P_{PL}(s = s_{ind} | u = E \rightarrow S) = 1/6$; see Table 6 for the respective values of other model components. The dependent Bayes net becomes more likely under a pragmatic interpretation only: by taking into account the fact that the speaker could have chosen a more informative utterance (e.g. “Sue goes on a skiing trip” (S)) to refer to the independent Bayes net (see Table 6), the listener learns that, most likely, the speaker believes in a connection between Sue going on a skiing trip and her passing the exam. Contrary to that, under a literal interpretation both Bayes nets remain equally likely as the chosen utterance, $E \rightarrow S$, is literally true in both states.

The predicted interpretation of Tom’s utterance, $E \rightarrow S$, with respect to the probability of the antecedent is shown in Figure 11(c): the listener’s degree of belief about the speaker’s beliefs in the antecedent remains at an expected value of 0.2, only the degree of belief related to the causal relation between E and S is influenced by the speaker’s utterance of the conditional. Crucially, since we so far only modelled the listener’s pragmatic inferences about the speaker’s beliefs, the listener’s observation ($C = c$) is not considered yet.

Literal meaning $\delta_{s \in [u]}$				Speaker $P_S(u s)$			
u	$E \rightarrow S$	S	likely S	u	$E \rightarrow S$	S	likely S
s_{dep}	1	0	0	s_{dep}	1	0	0
s_{ind}	1	1	1	s_{ind}	1/5	2/5	2/5

Literal listener $P_{\text{LL}}(s u)$				Pragmatic listener $P_{\text{PL}}(s u)$			
u	$E \rightarrow S$	S	likely S	u	$E \rightarrow S$	S	likely S
s_{dep}	1/2	0	0	s_{dep}	5/6	0	0
s_{ind}	1/2	1	1	s_{ind}	1/6	1	1

Table 6 Distributions for the Skiing Example with two Bayes nets s_{dep} , s_{ind} , utterances $\mathcal{U} = \{E \rightarrow S, S, \text{likely } S\}$ and $\alpha = 1$.

How should a listener change their beliefs in the light of a (pragmatically derived) belief about the speaker’s beliefs? That depends on the more general assumptions the listener makes about the speaker: is she trustworthy and well-informed on the subject matter at hand? Since nothing in the scenario described in Example (5) gives us reason to expect otherwise, we may follow the usual assumptions in Gricean belief-based reasoning that the speaker is cooperative and knowledgeable (e.g. [Geurts 2010](#)). Even if the precise effect of integrating beliefs of a cooperative and knowledgeable agent into one’s own beliefs are elusive, the effect in a scenario like the one at hand is most likely that the listener increases their own beliefs in the relation $E \rightsquigarrow^+ S$. To keep matters simple for a fully fleshed out numerical example, we just assume that the listener adopts exactly the same probabilistic beliefs as the inferred speaker beliefs.

This takes us to the last step of the explanation, where we look at the effect of the listener’s private knowledge that C is the case. Put differently, we look at the listener’s inference of the speaker’s beliefs about Sue passing the exam under the assumption that the speaker makes the same observation as the listener. Similar to [Douven 2012](#), we draw on world knowledge about the relation between C and S , where the former is highly unlikely when one does

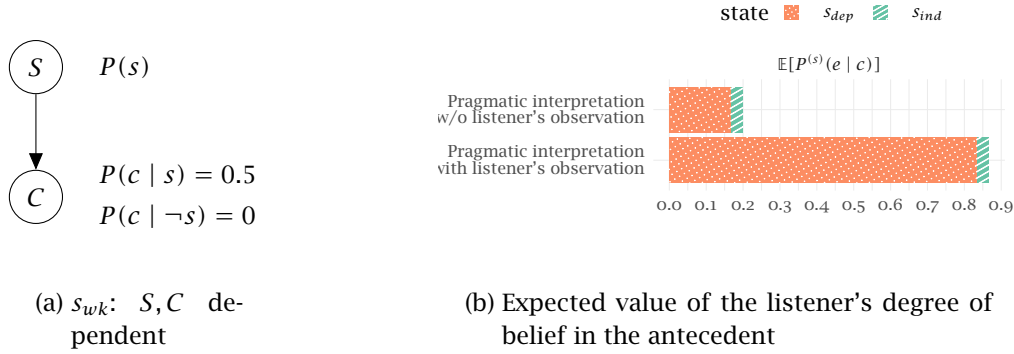


Figure 12 Bayes net for the assumed world knowledge in the Skiing Example (left) and results for the the listener's inference about the antecedent (right); E : pass exam, S : go skiing, C : buy skiing clothes.

not go on a skiing trip.²¹ This is reflected in the model by setting $P(c | \neg s)$ to 0 and $P(c | s)$ to 0.5 (see Figure 12(a)) as the context story does not provide any information concerning Sue's shopping behavior.

Based on this (world) knowledge about S and C (Bayes net s_{wk}), we can compute the listener's updated belief in the probability that Sue goes on a skiing trip given the listener's observation of Sue buying skiing clothes:

$$[11] \quad P^{s_{wk}}(s | c) = \frac{P^{s_{wk}}(c | s) \cdot P^{s_{wk}}(s)}{P^{s_{wk}}(c)} = \frac{0.5 \cdot P^{s_{wk}}(s)}{0.5 \cdot P^{s_{wk}}(s) + 0 \cdot P^{s_{wk}}(\neg s)} = 1$$

Assuming the updated probability of Sue going on a skiing trip, $P^{s_{wk}}(s | c)$ given in Equation [11], the updated probability of Sue passing the exam ($E = e$) for the two modeled states s_{dep} and s_{ind} is given in Equations [12] and [13]:

$$[12] \quad \begin{aligned} P^{(s_{dep})}(e | c) &= P^{(s_{dep})}(s, e | c) + P^{(s_{dep})}(\neg s, e | c) \\ &= P^{(s_{dep})}(e | s) \cdot P^{(s_{wk})}(s | c) + P^{(s_{dep})}(e | \neg s) \cdot P^{(s_{wk})}(\neg s | c) \\ &= 1 \cdot 1 + 0 \cdot 0 = 1 \end{aligned}$$

²¹ We argue that the relation between C and S is essential for the desired interpretation of the conditional. Imagine that Harry, the listener, had no idea what skiing is. In this case, both states should be modeled without the link between S and C making Harry's observation irrelevant with respect to his belief about Sue passing the exam (E). Therefore, his belief in the antecedent would remain unchanged. See Günther (2018) who also makes use of causal Bayes nets but argues that the intuitive interpretation of the Skiing example is independent of whether or not the relation between variables S and C is modeled.

A probabilistic model of communication with conditionals

$$\begin{aligned}
 [13] \quad P^{(s_{ind})}(e | c) &= P^{(s_{ind})}(s, e | c) + P^{(s_{ind})}(\neg s, e | c) \\
 &= P^{(s_{ind})}(e) \cdot P^{(s_{wk})}(s | c) + P^{(s_{ind})}(e) \cdot P^{(s_{wk})}(\neg s | c) \\
 &= 0.2 \cdot 1 + 0.2 \cdot 0 = 0.2
 \end{aligned}$$

These values emphasize the importance of taking into account the causal relations among variables: only for s_{dep} the degree of belief in E is influenced by the listener's observation of c . Given s_{ind} , the probability of Sue going on a skiing trip remains the same as without the listener's observation. The listener's updated belief in the antecedent, given the listener's independent observation of Sue buying skiing clothes, is then equal to the expected value of $P^{(s)}(e | c)$ ($s \sim P_{PL}(s | u = E \rightarrow S)$, $s \in [s_{dep}, s_{ind}]$) which is approximately 0.87, as spelled out in Equation [14]. Remember that $P_{PL}(s | u = E \rightarrow S)$ describes the listener's beliefs after the uptake of the conditional, that is, the listener is approximately 83% confident that the speaker's beliefs correspond to s_{dep} and approximately 17% confident that they correspond to s_{ind} . Figure 12(b) shows the predictions for the listener's posterior beliefs in the antecedent, with and without consideration of the listener's observation.

$$\begin{aligned}
 [14] \quad E[P^{(s)}(e | c)] &= \sum_{s_i \in \{s_{dep}, s_{ind}\}} P^{(s_i)}(e | c) \cdot P_{PL}(s_i | u = E \rightarrow S) = \\
 &P^{(s_{dep})}(e | c) \cdot P_{PL}(s_{dep} | u = E \rightarrow S) + \\
 &P^{(s_{ind})}(e | c) \cdot P_{PL}(s_{ind} | u = E \rightarrow S) = 1 \cdot 5/6 + 0.2 \cdot 1/6 \approx 0.87
 \end{aligned}$$

Conceptually, the expected value may be interpreted as integration of the listener's own observation with the information received from the speaker about the speaker's beliefs about the world that the listener takes over.

4.2 The Garden Party Example

Despite its structural similarity to the Skiing Example, for completeness, let us also briefly consider the Garden Party Example where, intuitively, the listener's degree of belief in the antecedent decreases.

The Garden Party Example (Douven 2012). Betty knows that Kevin, the son of her neighbors, was to take his driving test yesterday. She has no idea whether or not Kevin is a good driver; she deems it about as likely as not that Kevin passed the test. Betty notices that her neighbors have started to spade their garden. Then her mother, who is friends with Kevin's parents, calls her and tells her the following:

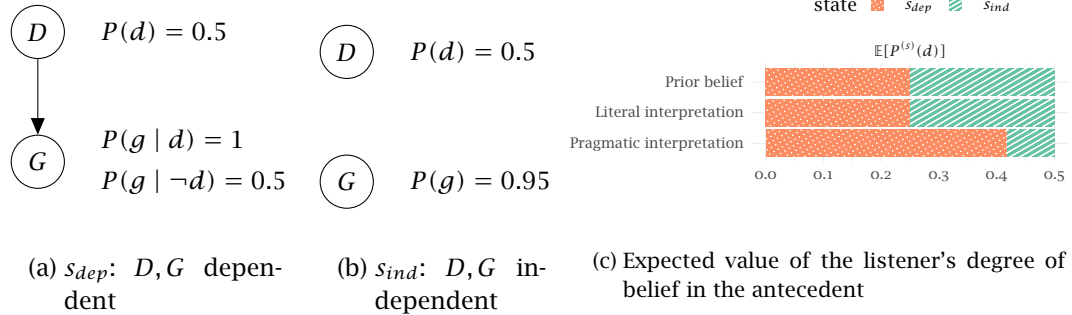


Figure 13 Bayes nets and results for Garden Party Example with $\mathcal{U} = \{G, D \rightarrow G, \text{likely } G\}$ and $\alpha = 3$, D : pass driving exam, G : throw garden party. Both states, s_{dep} and s_{ind} , are assigned equal prior probability.

- (6) If Kevin passed the driving test, his parents will throw a garden party.
 \rightsquigarrow listener belief in antecedent *decreases*

The difference between the Bayes nets shown in Figure 13 to the Bayes nets used in the Skiing example only lies in the intuitive instantiation of the prior probabilities. The results for the interpretation of the conditional in the Garden Party Example (Figure 13(c)) also falls in with the results in the Skiing Example: only under a pragmatic interpretation, the listener increases her belief in the Bayes net where both variables are connected (s_{dep}), but the listener's belief in the probability of the antecedent remains as prior to the speaker's utterance.

Again, to consider the listener's own beliefs, we assume that the listener simply takes over the speaker's beliefs communicated by the utterance of the conditional. The world knowledge that we draw on in this example, shown in Figure 14(a), concerns the fact that throwing a garden party is incompatible with spading the garden ($S = s$). Therefore, $P(s | g)$ is set to 0 and, due to the lack of more concrete information, $P(s | \neg g)$ is set to 0.5. Contrary to the Skiing Example, the listener's observation in this example is not evidence, but counterevidence for the antecedent: Betty observes her neighbors spading the garden, thus, the probability for a garden party decreases, $P^{(s_{wk})}(g | s) = 0$. While in s_{ind} , the degree of belief in D is not influenced by the truth or falsity of G , $P^{(s_{dep})}(d | \neg g) = 0$ (as opposed to $P^{(s_{dep})}(d | g) = 2/3$). The

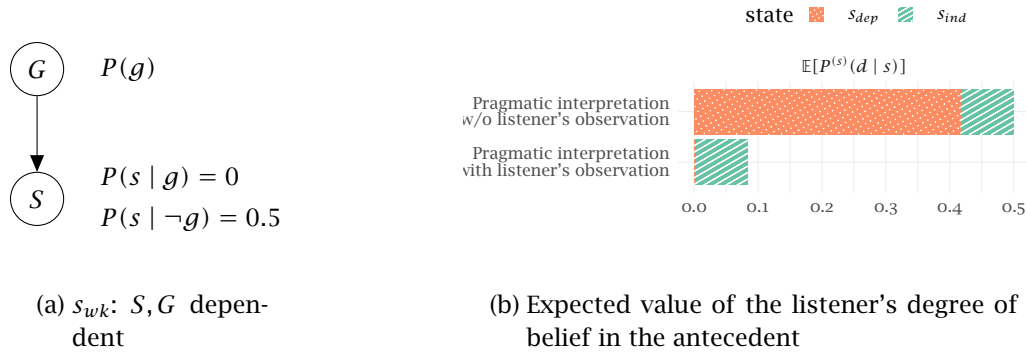


Figure 14 Bayes net for the assumed world knowledge in the Garden Party Example (left) and results for the the listener's inference about the antecedent (right); D : pass driving exam, G : throw garden party, S : spade garden.

combination of world knowledge about the connection between S and G , the speaker's utterance ($D \rightarrow G$), communicating a likely connection between D and G , and the listener's observation related to G therefore make Betty decrease her belief in the antecedent (Kevin passing the driving test), shown in Figure 14(b).

4.3 The Sundowners Example

The Sundowners Example is a case where, intuitively, the listener's degree of belief in the antecedent does not change much, if at all.

The Sundowners Example (Douven & Romeijn 2011). Sarah and Marian have arranged to go for sundowners at the Westcliff hotel tomorrow. Sarah feels there is some chance that it will rain, but thinks they can always enjoy the view from inside. To make sure, Marian consults the staff at the Westcliff hotel and finds out that in the event of rain, the inside area will be occupied by a wedding party. So she tells Sarah:


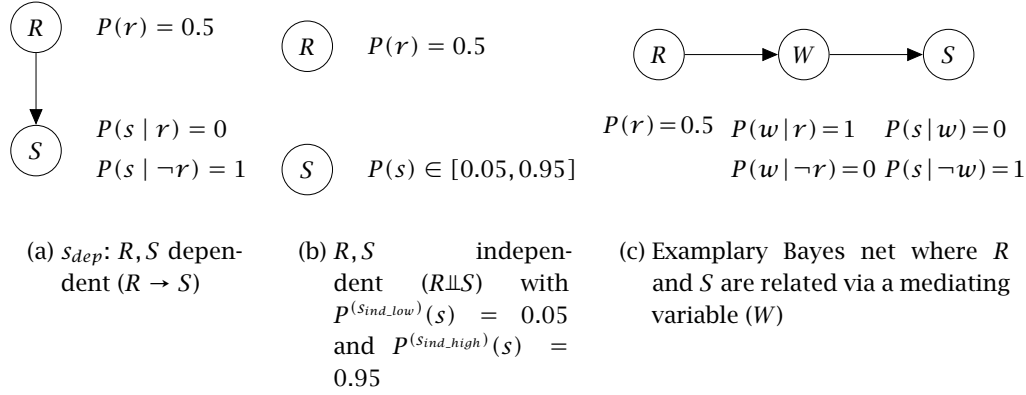
- (7) If it rains tomorrow, we cannot have sundowners at the Westcliff.
 \rightsquigarrow listener belief in antecedent remains *unchanged*

The intuition behind the formal treatment of the Sundowners Example, given below, is as follows. Even though the speaker’s utterance of the conditional does not seem infelicitous, it seems less natural than the conditionals in the previous two examples. This oddness is reflected in a response from the listener that is not far to seek: why should rain prevent them from having sundowners? Put differently, we expect the listener, who has a strong prior belief in the independence of R and S , to be somewhat surprised by the speaker’s utterance. The most rational explanation for the speaker’s utterance choice of the conditional $R \rightarrow \neg S$ is to give up the assumption of independence — at least if the integrity of the speaker is taken for granted. The listener’s surprise may be resolved by accommodating a third, latent variable which is neither observed nor part of the speaker’s utterance and acts as mediator between ‘rain’ and ‘having sundowners’. This would justify the speaker’s choice to utter the conditional $R \rightarrow \neg S$ which strongly suggests that (the speaker knows that) there is a (causal) connection between ‘rain’ and ‘having sundowners’. In other words, we propose that, as opposed to the previous examples, the speaker’s utterance in the Sundowners Example forces the listener to accommodate a mitigating variable, which forms a “causal bridge” between antecedent and consequent, so as to rationalize the speaker’s utterance. As nothing further is known about this newly introduced “causal bridge” or any other relevant variable, the result is that the listener’s beliefs about the variable ‘rain’ remains largely the same.

A formally explicit treatment of these ideas makes use of the Bayes nets shown in Figure 15(a)–15(c). As for the Skiing and the Garden Party Example, the speaker’s utterance of the conditional provokes a shift from the Bayes net where R and S are independent (Figure 15(b)) towards the Bayes net where they are causally related (Figure 15(a)). Yet, as can be seen in the leftmost panel of Figure 15(d), in this example, the listener has a strong prior belief in the former which explains the listener’s surprise resulting from the speaker’s utterance choice. Assuming independence of R and S , the listener would rather expect the speaker to choose a more informative utterance assertable than the conditional (e.g. , $\neg S$).

The RSA model of pragmatic language production entails a notion of *surprising utterances*. Given prior beliefs about states $P_{\text{prior}}(s)$ and the speaker’s assumed production probabilities $P_S(u | s)$, the pragmatic listener expects utterance u with probability:

$$P_{\text{PL}}(u) = \sum_s P_{\text{prior}}(s) P_S(u | s)$$



(d) Expected value of the listener's degree of belief in the causal relation between R and S (left), the probability of the antecedent (middle) and the joint event of antecedent and consequent (right), given the Bayes nets from Figure 15(a) and Figure 15(b).

Figure 15 Bayes nets and results for Sundowners Example with $\mathcal{U} = \{R \rightarrow \neg S, \text{likely } S, \text{likely } \neg S, S, \neg S\}$, $\alpha = 3$ and a prior probability of 0.85 for Bayes net s_{ind_high} , which is *a priori* most likely according to the context story; s_{ind_low} and s_{dep} are both assigned a prior probability of 0.075.

If the pragmatic listener only considers the three states represented in Figures 15(a) and 15(b), with a prior probability of 0.85 for s_{ind_high} and a prior probability of respectively 0.075 for s_{ind_low} and s_{dep} , the conditional $R \rightarrow \neg S$ is highly surprising in the sense that its expected occurrence probability is

very low (≈ 0.08 when $\alpha = 3$).²² This notion of listener surprise in the light of an unexpected utterance helps explain the intuition that an utterance of $R \rightarrow \neg S$ may trigger the desire to go look for an additional explanation which may rationalize the observed utterance.

One possibility of how the listener may rationalize a surprising utterance through a mitigating variable is shown in Figure 15(c) where W denotes the event of ‘a wedding taking place inside the hotel’ that represents any event which may possibly prevent the interlocutors from having sundowners at the hotel. Rather than denying the speaker’s integrity due to the speaker’s somewhat puzzling utterance, pragmatic reasoning, eminently causal-pragmatic reasoning, therefore allows the listener to infer a previously unexpected relation between R and S , that is able to explain the speaker’s utterance. Nevertheless, the conditional still seems to be incomplete as one could expect the speaker to be more informative about the reasons why the event of rain may prevent them from having sundowners.

Contrary to the Skiing and the Garden Party Example, in the Sundowners Example, the listener does not intuitively update her belief in the probability of the antecedent. Figure 15(d) (middle) shows that the listener’s predicted degree of belief in the antecedent remains at its prior value of 0.5 after processing the speaker’s utterance, assuming a literal or a pragmatic interpretation. The crucial difference here is that the listener does not make any further observations as in the previous examples. Yet, the speaker’s utterance of the conditional has a strong effect on the listener’s beliefs concerning the joint event of the antecedent and the consequent, shown in the right panel of Figure 15(d). Contrary to her prior belief that $P(r, s)$ is almost a matter of chance, the listener judges it almost impossible that the two events both hold at the same time upon receiving the conditional information $R \rightarrow \neg S$: in s_{dep} , R and S are mutually exclusive, thus $P(r, s) = 0$ and in s_{ind_low} where $R \rightarrow \neg S$ is assertable, $P(r, s)$ is close to 0.

In sum, the model is therefore able to account — by a single mechanism — for different interpretations with respect to the listener’s degree of belief in the antecedent, namely by the interplay of pragmatic reasoning and an adequate, explicit representation of the interlocutors’ prior probabilistic and in particular, *causal* beliefs.

²² $P_S(u = R \rightarrow \neg S \mid s_{ind_low}) \approx 0.06, P_S(u = R \rightarrow \neg S \mid s_{dep}) = 1$

5 On missing links & biscuits

Previous sections showed how the model of communication with conditionals presented here predicts that speakers use a conditional predominantly when there is a causal/inferential relation between antecedent and consequent and that therefore listeners will interpret conditionals accordingly. This brings up the obvious question as to how the advocated approach positions itself with regard to cases, prominently discussed in the literature, in which antecedent and consequent are clearly *not* causally or inferentially related. The absence of a relation between antecedent and consequent can either result in infelicity, as is the case in what we here call *missing-link conditionals* (Douven 2017), or trigger an altogether different kind of interpretation, as in what we here address as *biscuit conditionals* (e.g., Austin 1956, Geis & Lycan 1993). This section deals with each case in turn.

5.1 Missing-link conditionals

Missing-link conditionals have been used in support of the inferentialist position that the requirement of a causal/inferential connection between antecedent and consequent is a necessary condition for assertability of a conditional, arguably situated in the semantics of conditionals because it is unclear how else a pragmatic account could explain the *infelicity* of missing-link conditionals (e.g. Douven 2008, 2017, Krzyżanowska, Wenmackers & Douven 2014, Skovgaard-Olsen 2016). This position is concretely exemplified by the contrast pair in (8), given by Douven (2008).

- (8) There will be at least one heads in the first 1,000,000 tosses of this fair coin (h_{10^6}) *if*
- a. there is a heads in the first ten tosses. (h_{10})
 - b. * Chelsea wins the Champions League. (c)

What is remarkable about the contrast between the two sentences in (8) is this: on the one hand, both (8a) and (8b) arguably pass the minimal necessary requirement for assertability, namely that the probability of the consequent given the antecedent is high — in fact the difference between the relevant conditional probabilities is minute ($P(h_{10^6} | h_{10}) = 1$, $P(h_{10^6} | c) \approx 1$); on the other hand, while (8a) is intuitively acceptable, for instance in a context where the speaker wants to highlight the entailment relation to a listener who might otherwise not attend to it sufficiently, (8b) rather clearly is odd.

To account for the infelicity of (8b) and assertability of (8a), inferentialism stipulates that a conceivable inferential relation between antecedent and consequent is a necessary requirement for assertability, howsoever the inferential link may exactly be defined. It is not limited to deductive inferences, as modern inferentialism allows less strict inferential relations such as induction or abduction (e.g. Krzyżanowska, Wenmackers & Douven 2014, Douven 2017).²³ Inferentialists have criticized pragmatic explanations of the perceived infelicity of missing-link conditionals for remaining too vague about how the pragmatic processes may concretely look like (e.g., Douven 2017).

We argue here that (8b) is similar in kind to the Sundowners Example, yet more extreme, therefore leading to perceived infelicity. The Sundowners Example from the previous section showcases that there are contexts in which the utterance of a conditional is not questioned per se, but will neither be accepted without further ado. As there is no *obviously* conceivable connection between ‘rain’ and ‘not having sundowners at the hotel’, it seems quite natural for the listener to reply with a question asking for precisely this connection. We consider infelicitous missing-link conditionals like (8b) to be similar in kind but more extreme cases of the same variety: in (8b) the speaker provides so little information that the listener does not have enough cues to make sense of the conditional utterance from world knowledge alone. In other words, we maintain that the infelicity of (8b) is a result of a failure of the listener to see a connection between antecedent and consequent that could rationalize the speaker’s utterance choice.

To see how this is predicted by our model, we take the perspective of a listener, who tries to interpret a missing-link conditional and knows, from common sense world knowledge, that the antecedent and the consequent are two independent events. Given the assumption of independence, our hyper-rational speaker would never choose any of the conditional utterances (see Figure 6), and even for a speaker with a less optimal rationality parameter ($\alpha = 3$), we observe that conditional utterances only have an expected utterance choice probability $< 5\%$ when considering the speaker’s predictions for

²³ Several empirical studies (e.g., Krzyżanowska, Wenmackers & Douven 2013, Skovgaard-Olsen, Singmann & Klauer 2016, 2017) have shown that the link between antecedent and consequent has an influence on whether conditionals are accepted. Participants in a study from Douven & Verbrugge (2010) for instance interpreted conditionals differently depending on the type of the link (deductive, inductive or abductive). Yet, there is also empirical evidence supporting the view that the link between antecedent and consequent is attributable to discourse pragmatics rather than conventional semantics (e.g. see Cruz et al. 2016, Lassiter 2022).




Figure 16 Speaker’s expected utterance choice probabilities computed on speaker’s predictions, where $\alpha = 3$, for a set of 10,000 states sampled from the default context prior, grouped by utterance type and causal relation.

states with $r = A \perp\!\!\!\perp C$ from a set of 10,000 samples from the default context prior, as shown in Figure 16.

The reason for this low probability of choosing a conditional when A and C are assumed to be independent, is that very likely there will be an assertable, more informative utterance available; for the independent Bayes net, the joint probability tables where $A \rightarrow C$ is assertable (i.e., $P^{(s)}(c | a) \geq \theta$) also satisfy the assertability condition for a literal or a conjunction. Given the background knowledge of independence of the antecedent and the consequent and given that in these cases the speaker hardly ever chooses conditionals, the listener will naturally be highly surprised by the speaker’s utterance of a conditional. It is this notion of surprise on the part of the listener, based on the speaker’s expected utterance choice probabilities, that allows us to infer the infelicity of missing-link conditionals: due to the listener’s surprise, she will, arguably, want to look for another way of rationalizing the assertion. If no such option is forthcoming, the utterance feels infelicitous. Missing-link conditionals are therefore *not* accepted by the listener — contrary to the conditional in the Sundowners Example where the listener’s initial surprise can be repaired at the content-level since the listener is, after all, able to accommodate a connection between antecedent and consequent.

The infelicity of missing-link conditionals as we explain it here, may further be couched within the resource-rational, sampling-based approach to ex-

plain causal reasoning.²⁴ For example, Dasgupta, Schulz & Gershman (2017) show that cognitive biases like *subadditivity* or *superadditivity* can be explained by a comparison of people’s search for a plausible explanation of observed data with sampling algorithms, in particular MCMC sampling. They argue that due to cognitive load and time constraints, the number of samples that people can draw is restricted, such that the optimal answer is not found, giving rise to common cognitive biases. As regards to missing-link conditionals, the listener’s search for plausible Bayes nets with sufficient explanatory value to rationalize the speaker’s utterance, seems to fail (e.g., in Example 8b). Yet, when it comes to the interpretation of Biscuit conditionals (see Section 5.2 below) or to conditionals like in the Sundowners example, it seems reasonable for a listener to be able to quickly come up with more or less satisfying candidates.

5.2 Biscuit conditionals

Similar to missing-link conditionals, biscuit conditionals (BCs) are commonly considered a special kind of conditional as both lack the probably most characteristic feature of conditionals, the relation between antecedent and consequent. Unlike missing-link conditionals, biscuit conditionals are however felicitous, indeed quite common; see (9) for the classical example from Austin (1956).

(9) If you’re hungry, there are biscuits on the sideboard.

While Inferentialists have mostly excluded BCs from their analysis of conditionals altogether, we argue that the felicity of BCs can indeed be explained as a pragmatic phenomenon.²⁵ The model presented here helps explain why listeners, who will likely assume *a priori* that the consequent and the antecedent in example (9) are independent, would be surprised by an utterance of (9) if it were to be interpreted like a normal conditional. It is plausible that for some conditionals, like BCs, the listener will see a repair strategy which enables a rationalization of the conditional after all. However, this rationalization does not take place at the content-level by an attempt to find a relation between antecedent and consequent as (unsuccessfully) in missing-link conditionals or (successfully) in the Sundowners Example, but at a different

²⁴ Thanks to an anonymous reviewer for highlighting this interpretation.

²⁵ Recently, van Rooij & Schulz (2021a,b) proposed a generalization of their account for the assertability of conditionals which is able to account for BCs as well.

level, for instance at the level of speech-acts. Concretely, this idea could be integrated into our model by making the speaker's utterance choice additionally dependent on communicative goals, for example, whether the speaker wants to perform a speech-act (e.g., offer the listener some cookies), or just wants to make a plain assertion (e.g., inform the listener about the existence of biscuits on the sideboard).

Admittedly this sketch is just a first step towards a satisfying pragmatic account of BCs, as there is certainly more needed to fully explain their use and interpretation. We also do not, with emphasis, suggest that this kind of “surprise-repair” interpretation is actively entertained during each reception of a BC: BCs arguably provide sufficient secondary cues for the listener to trigger a BC-like interpretation, including intonation or, in some languages, word order (e.g., in German, where BCs can occur with verb-third (V3) verb order and not with the usual verb-second (V2) as seen in simple indicative conditionals).

5.3 Conditionals for the communication of *independence*

We assumed that, first and foremost, the speaker's aim is to communicate (probabilistic) beliefs about certain events and saw that, based on this assumption, the speaker should only choose to utter a conditional if no other, more informative utterance is assertable. However, in some circumstances, conditionals may be used in order to put emphasis on the fact that the consequent is *not* dependent on the antecedent. Consider the example given in (10):

- (10) If you study, you will pass and if you don't study, you will pass, so don't worry, you will pass!

Along the lines that we have been arguing, a conditional alike should be avoided since evidently the speaker could say the same thing with a more informative, shorter alternative, e.g., “*You will pass (no matter what)!*”. Notwithstanding, the conditional in (10) seems a natural utterance to say, in particular when the speaker wants to emphasize the relation, here the *independence* between studying and passing an exam. If we extended the set of alternative utterances available to our speaker by including this kind of combined conditionals ($A \rightarrow C \wedge \neg A \rightarrow C$), the listener would infer that the consequent holds true independent of the antecedent. We saw previously that our listener infers from the utterance of a conditional, that the speaker is likely uncertain

about the truth or falsity of the antecedent. For simplicity, let us assume the speaker refers to a state s where $P^{(s)}(a) = P^{(s)}(\neg a) = 0.5$. Since the speaker uttered $A \rightarrow C$, $P^{(s)}(c | a) \geq \theta$ must hold, and additionally, due to the speaker's utterance of the second conditional, $\neg A \rightarrow C$, $P^{(s)}(c | \neg a) \geq \theta$. Then, $P^{(s)}(c) = 0.5 \cdot P^{(s)}(c | a) + 0.5 \cdot P^{(s)}(c | \neg a) = 0.5 \cdot (P^{(s)}(c | a) + P^{(s)}(c | \neg a)) \geq \theta$. Therefore, in this state, it is also possible to assert the consequent straightaway. A speaker whose aim it is to communicate her uncertain beliefs should therefore prefer the simpler utterance, assuming higher cost for combined utterances like $A \rightarrow C \wedge \neg A \rightarrow C$. However, if the speaker's aim was two-minded, including the communication of her uncertain beliefs as well as highlighting the independence between the variables at hand, the combined conditional should become a likely utterance choice for the speaker: for almost all states in which both conditionals are assertable, $R = A \perp\!\!\!\perp C$ whereas C is also assertable in many states where $R = A \rightsquigarrow C$ or $R = C \rightsquigarrow A$.

6 Conclusion

We have taken a probabilistic approach to pragmatic reasoning about conditionals here. Our approach is a conservative one: we combined two independently motivated, well-established formalizations, namely the Rational-Speech-Act model from probabilistic pragmatics and causal Bayes nets, in order to predict a listener's posterior probabilistic beliefs about relevant events and how they are causally related. Our contribution is a formal, computational model of the use and interpretation of conditionals that may be a starting point for other models in this vein to gain further insights on how rich prior beliefs are updated in the light of conditionals.

We showed that our model vindicates a number of different pragmatic inferences observed in communication with conditionals. These include inferences such as the dependency between antecedent and consequent, the infelicity of missing-link conditionals or the tendency to interpret conditionals as biconditionals as in conditional perfection readings. When the model is supplied with particular prior beliefs representing concrete utterance contexts, it also makes predictions corresponding to what is considered their intuitive interpretation; we showed that the model helps to explain the listener's varying inferences with respect to the probability of the antecedent in three concrete utterance contexts from [Douven \(2012\)](#).

We explored the interpretation of conditionals under the assumption that conditionals are used to communicate stochastic information about co-

occurrence, while information about the underlying causal structure is communicated implicitly in our model. In some circumstances, the speaker's communicative goal might however go further than this and include the communication about causal information as well, bringing into play explicit causal language. An interesting extension of the model may investigate the relationship between conditionals, causality and causal language further by considering different alternative utterances, including explicit causal expressions that we did not address here.

The flexibility of the model we presented moreover allows future extensions to systematically investigate the factors influencing the interpretation of special types of conditionals such as biscuit conditionals. Other salient extensions of the present approach include applications to subjunctive conditionals, nested conditionals or conditionals with other logically complex antecedents or consequents. Further, the quantitative predictions of the model offer the possibility to test and improve it by a comparison with empirical data.

References

- Aloni, Maria. 2007. Expressing ignorance or indifference: Modal implicatures in bi-directional optimality theory. In Balder ten Cate & Henk Zeevat (eds.), *Logic, language and computation: Papers from the Tbilisi Symposium (TbiLLC)* (Lecture notes in computer science (LNCS) 4363), 1-20. Berlin: Springer. https://doi.org/10.1007/978-3-540-75144-1_1.
- Atlas, Jay David & Stephen C. Levinson. 1981. It-clefts, informativeness, and logical form. In Peter Cole (ed.), *Radical pragmatics*, 1-61. New York: Academic Press.
- Austin, J.L. 1956. Ifs and cans. *Proceedings of the British Academy* 42. 109-132. Reprinted as Austin 1979.
- Austin, J.L. 1979. Ifs and cans. In J. O. Urmson & Geoffrey J. Warnock (eds.), *Philosophical papers*, 3rd edn., 205-232. Cambridge: Cambridge University Press. <https://doi.org/10.1093/019283021X.003.0009>.
- van der Auwera, Johan. 1997. Pragmatics in the last quarter century: The case of conditional perfection. *Journal of Pragmatics* 27(3). 261-274. [https://doi.org/10.1016/s0378-2166\(96\)00058-6](https://doi.org/10.1016/s0378-2166(96)00058-6).
- Bennett, Jonathan F. 2003. *A philosophical guide to conditionals*. Oxford: Oxford University Press. <https://doi.org/10.1093/0199258872.001.0001>.

- Briggs, Ray. 2012. Interventionist counterfactuals. *Philosophical studies* 160(1). 139–166. <https://doi.org/10.1007/s11098-012-9908-5>.
- Brochhagen, Thomas, Michael Franke & Robert van Rooij. 2018. Coevolution of lexical meaning and pragmatic use. *Cognitive Science* 42(8). 2757–2789. <https://doi.org/10.1111/cogs.12681>.
- Burnett, Heather. 2019. Signalling games, sociolinguistic variation and the construction of style. *Linguistics and Philosophy* 42(5). 419–450. <https://doi.org/10.1007/s10988-018-9254-y>.
- Cheng, Patricia W. 1997. From covariation to causation: A causal power theory. *Psychological Review* 104(2). 367–405. <https://doi.org/10.1037//0033-295X.104.2.367>.
- Cruz, Nicole, David Over, Mike Oaksford & Jean Baratgin. 2016. Centering and the meaning of conditionals. *Proceedings of the Annual Meeting of the Cognitive Science Society* 38. 1104–1110.
- Dasgupta, Ishita, Eric Schulz & Samuel J. Gershman. 2017. Where do hypotheses come from? *Cognitive Psychology* 96. 1–25. <https://doi.org/10.1016/j.cogpsych.2017.05.001>.
- Degen, Judith, Robert D. Hawkins, Caroline Graf, Elisa Kreiss & Noah D. Goodman. 2020. When redundancy is useful: A bayesian approach to “overinformative” referring expressions. *Psychological Review* 127(4). 591–621. <https://doi.org/10.1037/rev0000186>.
- Díez, Francisco J. 1993. Parameter adjustment in Bayes networks: The generalized noisy OR-gate. In David Heckerman & Abe Mamdani (eds.), *Uncertainty in artificial intelligence (UAI): The ninth conference*, 99–105. San Mateo, CA: Morgan Kaufmann. <https://doi.org/10.1016/B978-1-4832-1451-1.50016-0>.
- Douven, Igor. 2008. The evidential support theory of conditionals. *Synthese* 164(1). 19–44. <https://doi.org/10.1007/s11229-007-9214-5>.
- Douven, Igor. 2012. Learning conditional information. *Mind and Language* 27(3). 239–263. <https://doi.org/10.1111/j.1468-0017.2012.01443.x>.
- Douven, Igor. 2017. How to account for the oddness of missing-link conditionals. *Synthese* 194. 1541–1554. <https://doi.org/10.1007/s11229-015-0756-7>.
- Douven, Igor, Shira Elqayam, Henrik Singmann & Janneke van Wijnbergen-Huitink. 2018. Conditionals and inferential connections: A hypothetical inferential theory. *Cognitive Psychology* 101. 50–81. <https://doi.org/10.1016/j.cogpsych.2017.09.002>.

- Douven, Igor & Jan Willem Romeijn. 2011. A new resolution of the Judy Benjamin problem. *Mind* 120(479). 637-670. <https://doi.org/10.1093/mind/fzr051>.
- Douven, Igor & Sara Verbrugge. 2010. The Adams family. *Cognition* 117(3). 302-318. <https://doi.org/10.1016/j.cognition.2010.08.015>.
- Douven, Igor & Sara Verbrugge. 2012. Indicatives, concessives, and evidential support. *Thinking & Reasoning* 18(4). 480-499. <https://doi.org/10.1080/13546783.2012.716009>.
- Earman, John. 1992. *Bayes or bust? A critical examination of Bayesian confirmation theory*. Cambridge: MIT Press.
- Edgington, Dorothy. 1995. On conditionals. *Mind* 104(414). 235-329. <https://doi.org/10.1093/mind/104.414.235>.
- Eva, Benjamin, Stephan Hartmann & Soroush Rafiee Rad. 2019. Learning from conditionals. *Mind* 129(514). 461-508. <https://doi.org/10.1093/mind/fzz025>.
- Evans, Johathan St. B. T. & David E. Over. 2004. *If*. Oxford: Oxford University Press.
- Evans, Jonathan St. B. T., Simon J. Handley, Helen Neilens & David E. Over. 2007. Thinking about conditionals: A study of individual differences. *Memory & Cognition* 35(7). 1772-1784. <https://doi.org/10.3758/BF03193509>.
- Fernbach, Philip M. & Adam Darlow. 2010. Neglect of alternative causes in predictive but not diagnostic reasoning. *Psychological Science* 21(3). 329-336. <https://doi.org/10.1177/0956797610361430>.
- Fernbach, Philip M., Adam Darlow & Steven A. Sloman. 2011. Asymmetries in predictive and diagnostic reasoning. *Journal of experimental psychology: General* 140(2). 168-185. <https://doi.org/10.1037/a0022100>.
- Fernbach, Philip M. & Bob Rehder. 2013. Cognitive shortcuts in causal inference. *Argument and Computation* 4(1). 64-88. <https://doi.org/10.1080/19462166.2012.682655>.
- de Finetti, Bruno. 1936. La logique de la probabilité. In *Actes du congrès international de philosophie scientifique*, vol. 4, 565-573. Paris: Hermann Editeurs. <http://www.brunodefinetti.it/Opere/logiquedelaProbabilite.pdf>.
- von Fintel, Kai. 2001. Conditional strengthening: A study in implicature. Ms., MIT. <https://web.mit.edu/fintel/fintel-2001-condstrength.pdf>.
- van Fraassen, Bas C. 1976. Probabilities of conditionals. In William L. Harper & Cliff A. Hooker (eds.), *Foundations of probability theory, statistical inference, and statistical theories of science* (Foundations and Philosophy of

- Epistemic Applications of Probability Theory 1), 261–308. Dordrecht: Reidel. https://doi.org/10.1007/978-94-010-1853-1_10.
- Frank, Michael C. & Noah D. Goodman. 2012. Predicting pragmatic reasoning in language games. *Science* 336(6084). 998. <https://doi.org/10.1126/science.1218633>.
- Franke, Michael. 2011. Quantity implicatures, exhaustive interpretation, and rational conversation. *Semantics & Pragmatics* 4(1). 1–82. <https://doi.org/10.3765/sp.4.1>.
- Franke, Michael. 2014. Pragmatic reasoning about unawareness. *Erkenntnis* 79(4). 729–767. <https://doi.org/10.1007/s10670-013-9464-1>.
- Franke, Michael & Leon Bergen. 2020. Theory-driven statistical modeling for semantics and pragmatics: A case study on grammatically generated implicature readings. *Language* 96(2). e77–e96. <https://doi.org/10.1353/lan.2020.0034>.
- Franke, Michael & Gerhard Jäger. 2016. Probabilistic pragmatics, or why Bayes' rule is probably important for pragmatics. *Zeitschrift für Sprachwissenschaft* 35(1). 3–44. <https://doi.org/10.1515/zfs-2016-0002>.
- Franke, Michael & Tikitou de Jager. 2011. Now that you mention it: Awareness dynamics in discourse and decisions. In Anton Benz, Christian Ebert, Gerhard Jäger & Robert van Rooij (eds.), *Language, games, and evolution: Trends in current research on language and game theory* (Lecture notes in computer science (LNAI) 6207), 60–91. Heidelberg: Springer. https://doi.org/10.1007/978-3-642-18006-4_4.
- Gates, Monica A., Tess Veuthey, Michael Henry Tessler, Kevin A. Smith, Tobias Gerstenberg, Laurie Bayet & Josh Tenenbaum. 2018. Tiptoeing around it: Inference from absence in potentially offensive speech. *Proceedings of the Annual Meeting of the Cognitive Science Society* 40. 1693–1698. <https://cicl.stanford.edu/papers/gates2018tiptoeing.pdf>.
- Geis, Michael L & Arnold M Zwicky. 1971. On invited inferences. *Linguistic Inquiry* 2(4). 561–566. <http://www.jstor.org/stable/4177664>.
- Geis, Michael L. & William G. Lycan. 1993. Nonconditional conditionals. *Philosophical Topics* 21(2). 35–56. <http://www.jstor.org/stable/43154153>.
- Geurts, Bart. 2010. *Quantity implicatures*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511975158>.
- Goodman, Noah D. & Michael C. Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences* 20(11). 818–829. <https://doi.org/10.1016/j.tics.2016.08.005>.

- Goodman, Noah D. & Andreas Stuhlmüller. 2013. Knowledge and implicature: Modeling language understanding as social cognition. *Topics in Cognitive Science* 5(1). 173–184. <https://doi.org/10.1111/tops.12007>.
- Goodman, Noah D. & Andreas Stuhlmüller. 2014. *The design and implementation of probabilistic programming languages*. <http://dippl.org>. Last accessed 2020-09-07.
- Grice, H. Paul. 1989. Indicative conditionals. In *Studies in the way of words*. Harvard University Press.
- Grice, H. Paul. 2013. Logic and conversation. In Peter Cole & Jerry L. Morgan (eds.), *Speech acts* (Syntax and Semantics 3), 41–58. Leiden: Brill. https://doi.org/10.1163/9789004368811_003.
- Griffiths, Thomas L. & Joshua B. Tenenbaum. 2005. Structure and strength in causal induction. *Cognitive psychology* 51(4). 334–384. <https://doi.org/10.1016/j.cogpsych.2005.05.004>.
- Grove, Adam J. & Joseph Y. Halpern. 1997. Probability update: Conditioning vs. cross-entropy. In *Uncertainty in artificial intelligence (UAI): The thirteenth conference*, 208–214. <https://doi.org/10.5555/2074226.2074251>.
- Günther, Mario. 2018. Learning conditional and causal information by Jeffrey imaging on Stalnaker conditionals. *Journal of Philosophical Logic* 47(4). 851–876. <https://doi.org/10.1007/s10992-017-9452-z>.
- Hadjichristidis, Constantinos, Rosemary J. Stevenson, David E. Over, Steven A. Sloman, Jonathan St. B. T. Evans & Aidan Feeney. 2001. On the evaluation of ‘if p then q’ conditionals. *Proceedings of the Annual Meeting of the Cognitive Science Society* 21. 409–414. <https://escholarship.org/uc/item/87f7c704>.
- Hagmayer, York & Michael R. Waldmann. 2007. Inferences about unobserved causes in human contingency learning. *The Quarterly Journal of Experimental Psychology* 60(3). 330–355. <https://doi.org/10.1080/17470210601002470>.
- Hájek, Alan. 1989. Probabilities of conditionals – Revisited. *Journal of Philosophical Logic* 18(4). 423–428. <https://doi.org/10.1007/BF00262944>.
- Heifetz, Aviad, Martin Meier & Burkhard C. Schipper. 2006. Interactive unawareness. *Journal of Economic Theory* 130. 78–94. <https://doi.org/10.1016/j.jet.2005.02.007>.
- Herbstritt, Michele & Michael Franke. 2019. Complex probability expressions & higher-order uncertainty: Compositional semantics, probabilistic pragmatics & experimental data. *Cognition* 186. 50–71. <https://doi.org/10.1016/j.cognition.2018.11.013>.

- Hiddleston, Eric. 2005. A causal theory of counterfactuals. *Noûs* 39(4). 632–657. <https://doi.org/10.1111/j.0029-4624.2005.00542.x>.
- Horn, Laurence R. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Deborah Schiffrin (ed.), *Meaning, form, and use in context: Linguistic applications* (Georgetown University Round Table on Languages and Linguistics (GURT)), 11–42. Washington, D.C.: Georgetown University Press.
- Horn, Laurence R. 2000. From if to iff: Conditional perfection as pragmatic strengthening. *Journal of Pragmatics* 32(3). 289–326. [https://doi.org/https://doi.org/10.1016/S0378-2166\(99\)00053-3](https://doi.org/https://doi.org/10.1016/S0378-2166(99)00053-3).
- Hytinen, Antti, Frederick Eberhardt & Patrik O. Hoyer. 2011. Noisy-OR Models with Latent Confounding. *Uncertainty in Artificial Intelligence (UAI)* 27. 363–372. <https://doi.org/10.48550/arXiv.1202.3735>.
- Icard, Thomas F. & Noah D. Goodman. 2015. A resource-rational approach to the causal frame problem. *Proceedings of the Annual Meeting of the Cognitive Science Society* 37. 962–967. <http://cocolab.stanford.edu/papers/IcardGoodman2015-Cogsci.pdf>.
- Kao, Justine T., Leon Bergen & Noah D. Goodman. 2014. Formalizing the pragmatics of metaphor understanding. *Proceedings of the Annual Meeting of the Cognitive Science Society* 36. 719–724. <https://escholarship.org/uc/item/09h3p4cz>.
- Kao, Justine T., Jean Y. Wu, Leon Bergen & Noah D. Goodman. 2014. Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences (PNAS)* 111(33). 12002–12007. <https://doi.org/10.1073/pnas.1407479111>.
- Katzir, Roni. 2007. Structurally-defined alternatives. *Linguistics and Philosophy* 30. 669–690. <https://doi.org/10.1007/s10988-008-9029-y>.
- Kaufmann, Stefan. 2004. Conditioning against the grain. *Journal of Philosophical Logic* 33(6). 583–606. <https://doi.org/10.1023/B:LOGI.0000046142.51136.bf>.
- Kaufmann, Stefan. 2013. Causal premise semantics. *Cognitive Science* 37(6). 1136–1170. <https://doi.org/10.1111/cogs.12063>.
- Khoo, Justin. 2016. Probabilities of conditionals in context. *Linguistics and Philosophy* 39(1). 1–43. <https://doi.org/10.1007/s10988-015-9182-z>.
- Kratzer, Angelika. 1991. Modality. In Arnim von Stechow & Dieter Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung*, 639–650. Berlin: Walter de Gruyter. <https://doi.org/10.1515/9783110126969.7.639>.

- Krynski, Tevye R. & Joshua B. Tenenbaum. 2007. The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General* 136(3). 430–450. <https://doi.org/10.1037/0096-3445.136.3.430>.
- Krzyżanowska, Karolina, Sylvia Wenmackers & Igor Douven. 2013. Inferential conditionals and evidentiality. *Journal of Logic, Language and Information* 22(3). 315–334. <https://doi.org/10.1007/s10849-013-9178-4>.
- Krzyżanowska, Karolina, Sylvia Wenmackers & Igor Douven. 2014. Rethinking Gibbard's riverboat argument. *Studia Logica* 102(4). 771–792. <https://doi.org/10.1007/s11225-013-9507-2>.
- Lassiter, Daniel. 2017. Probabilistic language in indicative and counterfactual conditionals. *Semantics and Linguistic Theory (SALT)* 27. 525–546. <https://doi.org/10.3765/salt.v27i0.4188>.
- Lassiter, Daniel. 2018. Talking about (quasi-)higher-order uncertainty. In Cleo Condoravdi & Tracy Holloway King (eds.), *Tokens of meaning: Papers in honor of Lauri Karttunen*. Stanford, CA: CSLI Publications.
- Lassiter, Daniel. 2020. What we can learn from how trivalent conditionals avoid triviality. *Inquiry* 63(9-10). 1087–1114. <https://doi.org/10.1080/0020174X.2019.1698457>.
- Lassiter, Daniel. 2022. Decomposing relevance in conditionals. *Mind and Language*. 644–668. <https://doi.org/10.1111/mila.12418>.
- Lassiter, Daniel & Noah D. Goodman. 2017. Adjectival vagueness in a Bayesian model of interpretation. *Synthese* 194(10). 3801–3836. <https://doi.org/10.1007/s11229-015-0786-1>.
- Levinson, Stephen C. 2000. *Presumptive meanings: The theory of generalized conversational implicature*. Cambridge: MIT Press. <https://doi.org/10.7551/mitpress/5526.001.0001>.
- Lewis, David. 1976. Probabilities of conditionals and conditional probabilities. *Philosophical Review* 85(3). 297–315. <https://doi.org/10.2307/2184045>.
- Lucas, Christopher G. & Charles Kemp. 2015. An improved probabilistic account of counterfactual reasoning. *Psychological Review* 122(4). 700–734. <https://doi.org/10.1037/a0039655>.
- Matsumoto, Yo. 1995. The conversational condition on Horn scales. *Linguistics and Philosophy* 18(1). 21–60. <https://doi.org/10.1007/BF00984960>.
- Milne, Peter. 1997. Bruno de Finetti and the logic of conditional events. *The British Journal for the Philosophy of Science* 48(2). 195–232. <https://doi.org/10.1093/bjps/48.2.195>.

- Moldovan, Andrei. 2013. Denying the antecedent and conditional perfection again. In Dima Mohammed & Marcin Lewiński (eds.), *Virtues of argumentation: Proceedings of the 10th international conference of the Ontario Society for the Study of Argumentation (OSSA)*, 1–10. Windsor, ON: OSSA. <https://scholar.uwindsor.ca/ossaarchive/OSSA10/papersandcommentaries/117>.
- Moss, Sarah. 2015. On the semantics and pragmatics of epistemic vocabulary. *Semantics and Pragmatics* 8. 1–81. <https://doi.org/10.3765/sp.8.5>.
- Newstead, Stephen E. 1997. Conditional reasoning with realistic material. *Thinking & Reasoning* 3(1). 49–76. <https://doi.org/10.1080/135467897394428>.
- Oberauer, Klaus, Andrea Weidenfeld & Katrin Fischer. 2007. What makes us believe a conditional? The roles of covariation and causality. *Thinking & Reasoning* 13(4). 340–369. <https://doi.org/10.1080/13546780601035794>.
- Oberauer, Klaus & Oliver Wilhelm. 2003. The meaning(s) of conditionals: Conditional probabilities, mental models, and personal utilities. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29(4). 680–693. <https://doi.org/10.1037/0278-7393.29.4.680>.
- Pearl, Judea. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Burlington, MA: Morgan Kaufmann. <https://doi.org/10.1016/C2009-0-27609-4>.
- Pearl, Judea. 2009. *Causality*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511803161>.
- Pearl, Judea. 2013. Structural counterfactuals: A brief introduction. *Cognitive Science* 37(6). 977–985. <https://doi.org/10.1111/cogs.12065>.
- Qing, Ciyang & Michael Franke. 2015. Variations on a bayesian theme: Comparing bayesian models of referential reasoning. In Henk Zeevat & Hans-Christian Schmitz (eds.), *Bayesian natural language semantics and pragmatics* (Language, Cognition, and Mind (LCAM) 2), 201–220. Cham: Springer. https://doi.org/10.1007/978-3-319-17064-0_9.
- Qing, Ciyang, Noah D. Goodman & Daniel Lassiter. 2016. A rational speech-act model of projective content. *Proceedings of the Annual Meeting of the Cognitive Science Society* 38. 1110–1115. <https://cocolab.stanford.edu/papers/QingGoodmanLassiter2016-Cogsci.pdf>.
- Quine, Willard Van Orman. 1965. *Elementary logic: Revised edition*. New York: Harper.

- Rips, Lance J. 2010. Two causal theories of counterfactual conditionals. *Cognitive Science* 34(2). 175–221. <https://doi.org/10.1111/j.1551-6709.2009.01080.x>.
- Roberts, Craige. 1996. Information structure in discourse: Towards an integrated formal theory of pragmatics. In Jae Hak Yoon & Andreas Kathol (eds.), *Papers in semantics* (Working Papers in Linguistics 49). The Ohio State University, Department of Linguistics. A revised version was published in Roberts 2012. https://kb.osu.edu/dspace/bitstream/handle/1811/81500/WPL_49_Summer_1996_091.pdf.
- Roberts, Craige. 2012. Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics* 5(1). 1–69. <https://doi.org/10.3765/sp.5.6>.
- van Rooij, Robert & Katrin Schulz. 2019. Conditionals, causality and conditional probability. *Journal of Logic, Language and Information* 28(1). 55–71. <https://doi.org/10.1007/s10849-018-9275-5>.
- van Rooij, Robert & Katrin Schulz. 2021a. Conditionals as representative inferences. *Axiomathes* 31(3). 437–452. <https://doi.org/10.1007/s10516-020-09477-9>.
- van Rooij, Robert & Katrin Schulz. 2021b. Why those biscuits are relevant and on the sideboard. *Theoria* 87(3). 704–712. <https://doi.org/10.1111/theo.12309>.
- Rothschild, Daniel. 2015. Conditionals and propositions in semantics. *Journal of Philosophical Logic* 44(6). 781–791. <https://doi.org/10.1007/s10992-015-9359-5>.
- Santorio, Paolo. 2016. Interventions in premise semantics. *Philosophers' Imprint* 19(1). 1–27. <https://doi.org/2027/spo.3521354.0019.001>.
- Schuster, Sebastian & Judith Degen. 2019. Speaker-specific adaptation to variable use of uncertainty expressions. *Proceedings of the Annual Meeting of the Cognitive Science Society* 41. 2769–2775. <https://alpslab.stanford.edu/papers/2019SchusterDegen.pdf>.
- Scontras, Gregory, Micahel Henry Tessler & Michael Franke. 2017. *Probabilistic language understanding: An introduction to the Rational Speech Act framework*. <https://www.problang.org>. Retrieved 2021-3-24.
- Singmann, Henrik, Karl Christoph Klauer & David Over. 2014. New normative standards of conditional reasoning and the dual-source model. *Frontiers in Psychology* 5. <https://doi.org/10.3389/fpsyg.2014.00316>.

- Skovgaard-Olsen, Niels. 2016. Motivating the relevance approach to conditionals. *Mind and Language* 31(5). 555–579. <https://doi.org/10.1111/mila.12120>.
- Skovgaard-Olsen, Niels, Henrik Singmann & Karl Christoph Klauer. 2016. The relevance effect and conditionals. *Cognition* 150. 26–36. <https://doi.org/10.1016/j.cognition.2015.12.017>.
- Skovgaard-Olsen, Niels, Henrik Singmann & Karl Christoph Klauer. 2017. Relevance and reason relations. *Cognitive Science* 41. 1202–1215. <https://doi.org/10.1111/cogs.12462>.
- Stalnaker, Robert. 1968. A theory of conditionals. In Nicholas Rescher (ed.), *Studies in logical theory*, 98–112. Oxford: Blackwell. Reprinted as *Stalnaker 1981*.
- Stalnaker, Robert. 1981. A theory of conditionals. In William L. Harper, Robert Stalnaker & Glenn Pearce (eds.), *Ifs* (The University of Western Ontario Series in Philosophy of Science 15), 41–55. Dordrecht: Springer. https://doi.org/10.1007/978-94-009-9117-0_2.
- Stalnaker, Robert & Richard Jeffrey. 1994. Conditionals as random variables. In Ellery Eells & Brian Skyrms (eds.), *Probability and conditionals: Belief revision and rational decision*, 31–46. Cambridge: Cambridge University Press.
- Swanson, Eric. 2015. The application of constraint semantics to the language of subjective uncertainty. *Journal of Philosophical Logic* 45. 1–26. <https://doi.org/10.1007/s10992-015-9367-5>.
- Tenenbaum, Joshua B. & Thomas L. Griffiths. 2003. Theory-based causal inference. *Advances in neural information processing systems (NIPPS)* 15. 43–50. <http://papers.neurips.cc/paper/2332-theory-based-causal-inference.pdf>.
- Tenenbaum, Joshua B., Charles Kemp, Tom L. Griffiths & Noah D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331(6022). 1279–1285. <https://doi.org/10.1126/science.1192788>.
- Ülkümen, Gülden, Craig R Fox & Bertram F Malle. 2016. Two dimensions of subjective uncertainty: Clues from natural language. *Journal of Experimental Psychology: General* 145(10). 1280–1297. <https://doi.org/10.1037/xge0000202>.
- Vandenburgh, Jonathan. 2020. Conditional learning through causal models. *Synthese* 199. 2415–2437. <https://doi.org/10.1007/s11229-020-02891-x>.
- Yalcin, Seth. 2012. Bayesian expressivism. *Proceedings of the Aristotelian Society* 112(2). 123–160. <https://doi.org/10.1111/j.1467-9264.2012.00329.x>.

A probabilistic model of communication with conditionals

Yoon, Erica J., Michael Henry Tessler, Noah D. Goodman & Michael C. Frank. 2016. Talking with tact: Polite language as a balance between kindness and informativity. *Proceedings of the Annual Meeting of the Cognitive Science Society* 38. 2771-2776. https://langcog.stanford.edu/papers_new/yoon-2016-cogsci.pdf.

Britta Grusdt
Wachsbleiche 27
49090 Osnabrück, Germany
britta.grusdt@uni-osnabrueck.de

Daniel Lassiter
Dugald Stewart Building
3 Charles Street, Edinburgh
EH8 9AD, United Kingdom
dlassite@ed.ac.uk

Michael Franke
Geschwister-Scholl-Platz
72074 Tübingen, Germany
michael.franke@uni-tuebingen.de