Independence and Decision-Contexts for Non-Interference Conditionals*

Michael Franke
Institute for Logic, Language and Computation
Universiteit van Amsterdam

m.franke@uva.nl

Abstract

Non-interference conditionals are conditional sentences whose antecedent and consequent are not conditionally related, e.g.: "If you are hungry, there are biscuits on the shelf." Previous accounts struggled to explain two phenomena in connection with these sentences. Firstly, non-interference conditionals seem to imply their consequents. This can be explained, if we evaluate conditionals with respect to the epistemic state of a speaker and spell out what it means to say that a speaker believes in the conditional independence of propositions. Secondly, non-interference conditionals seem to be conditional assertions. This, too, can be explained even with a standard semantics for conditionals, if we employ a rich context-model, i.e. evaluate propositional information against the background of a decision problem of the hearer.

1 Non-Interference Conditionals

During one of his recorded shows the American comedian Demetri Martin told the following joke, much to the amusement of his audience:

> She was amazing. I never met a woman like this before. She showed me to the dressing room. She said: "If you need anything, I'm Jill." I was like: "Oh, my God! I never met a woman before with a *conditional identity*." [Laughter] "What if I don't need anything? Who are you?" — "If you don't need anything, I'm Eugene." [More laughter]
> (Demetri Martin, *These are jokes*)

Martin's joke is possible because of a peculiarity of certain conditional sentences. Some conditional sentences relate propositions that have no conditional relationship. This is by no means contradictory or paradoxical. The sentence

(1) If you need anything, I'm Jill.

links the clauses "you need anything" and "I'm Jill" in a conditional construction, but semantically we may naturally perceive the propositions expressed by these clauses as conditionally unrelated; the name of the woman does not depend on whether the addressee needs anything or not. To humorously misapprehend such conditional sentences, as Martin does, is to pretend to see a conditional relationship where none exists.

*Thanks to Robert van Rooij, Sven Lauer, Anton Benz and Tikitu de Jager for help and discussion.

Examples that would lend themselves to similar joking have been discussed as a special case of conditionals from a variety of angles under a variety of names (see Siegel (2006) for a recent treatment and more references). I will speak of *non-interference conditionals* (NICs), because of the intuitive conditional independence of antecedent and consequent in such sentences. However, not all sentences that are NICs in this sense are exactly alike. A reasonable distinction is given by Günthner (1999) based on a corpus study of spoken German: she lists meta-communicative conditionals (2a) and discourse-structuring conditionals (2b), next to relevance conditionals like (1).

(2) a. If I'm honest, I actually like Heidegger.
b. If we now turn to the last point of order, the fund cuts have been tremendous.

Meta-communicative conditionals like (2a) have antecedents which relate in some fashion to communicative rules or the actual linguistic conduct of the speaker. All sorts of politeness hedging would file under this heading. Witness, for instance, phrases like "if I may say so", "if you ask me" or "if I may interrupt". Discourse-structuring conditionals like (2b) seem to introduce or change a topic. These are, perhaps, the least likely variety to be found in written language. If spoken, they show a characteristic pause after the *if*-clause, which nearly exclusively occurs sentence-initially. Relevance conditionals like (1) and (3) are constructions whose antecedent intuitively gives a sufficient — and often also necessary — condition for the *relevance* of the information given in the consequent.[1]

(3) a. If you are are hungry, there are biscuits on the shelf.
b. If you want some, there are biscuits on the shelf.

It is clear that NICs are quite unlike ordinary conditionals such as (4).

(4) a. If it does not rain, we will eat outside.
b. If the butler has not killed the baroness, the gardener has.

Still, NICs are very natural things to say; such sentences are frequent in spontaneous speech and this for good reasons. One consequence of this is that it is easy to imagine a lively context of utterance with rather distinct properties when confronted with naked examples of NICs even in the laboratory context of a linguistic paper such as this. Consequently, the main idea behind this treatment of NICs is to explain some of the characteristic differences between NICs and ordinary conditionals *pragmatically*, and that is to say not in terms of, e.g., an assumed grammatical difference between kinds of conditionals, but rather in terms of (sharpened formalizations of) our rich and solid intuitions about the contexts in which NICs would naturally be used. The thesis to be advanced and defended here is that NICs can be given a standard semantic analysis, yet still their elaborate contextual meanings can

[1]Strictly speaking, two more kinds of sentences are NICs in the given sense: for one, there are so-called monkey's-uncle conditionals like (i) with an obviously false consequent as an emphatic way to deny or reject the antecedent; for another, there are concessive conditionals with an 'even-if' reading like (ii).

(i) If that's true, I'm a monkey's uncle.

(ii) That match is wet. If you strike it, it won't light.

For the present discussion, both of these are of no interest. I exclude them when I speak of NICs.

be explained. We just have to make proper use of our understanding of the contexts in which NICs would be used.

What are then the differences between NICs and ordinary conditionals? One frequently discussed difference is this: NICs seem to imply the truth of their consequent proposition. For instance, any relevance conditional in (3) somehow conveys that there are indeed biscuits on the shelf, no matter whether this information is relevant or not. Of course, this is different for ordinary conditionals like (4) and therefore a theory of NICs should explain how the feeling of entailment of the consequent comes about. The main aim of section 2 is to account for this perceived entailment relation.

Another peculiar meaning intuition to be taken seriously and explained is the following. Somehow NICs appear very much like *conditional assertions*: not the truth of the consequent but the assertion thereof depends on the antecedent. If the antecedent is false, then, intuitively, the speech-act associated with the main clause is somehow not feasible: for meta-communicative conditionals (2a) some politeness rule or felicity condition would be infringed; for discourse-structuring conditionals (2b) the assertion would be off-topic; for relevance conditionals (1) and (3) the assertion would be irrelevant. The raison d'être of NICs is therefore, the intuition continues, to protect speakers from acting inappropriately; speakers perform speech-acts only conditionally, if they are uncertain whether these acts are feasible.

Many theories of NICs have taken this intuition at face value. In order to explain how a conditional sentence like (3a), for instance, can be interpreted as a conditional assertion, linguists and philosophers have variously advanced theories that, in crude outline, either postulate an elliptical performative (5a) or some abstract illocutionary force operator (5b).[2]

(5) a. If you are hungry, (I hereby say to you that) there are biscuits on the shelf.

b. If you are hungry, ASSERT("there are biscuits on the shelf").

Even when we neglect the intricacies of individual proposals, their merits and flaws, it is safe to say what is unappealing about any such account. For one, conditional speech-acts, if taken seriously, are very peculiar entities — where else in life do you perform your actions conditionally? — whose properties can only be assessed via exactly those sentences' meanings whose meaning they are to explain. For another, any account that strives to analyze NICs as conditional assertions along the lines of (5), even if that account is more carefully elaborated as in Siegel (2006), postulates hidden linguistic structure of some kind, be that performatives, speech-act operators or else. Clearly, where possible, a simple standard semantics of conditional sentences would be preferred. I argue in section 3 that we can treat NICs as standard conditionals and still account for the intuition that assertions are somehow conditionalized, as long as we assume a rich enough notion of context.

[2] A particularly outspoken proposal of the performative analysis in (5a) is van der Auwera (1986). The perhaps most extreme version of the conditional speech-act hypothesis is the theory proposed by DeRose and Grandy (1999) who argue that *all* conditionals are such conditional assertions.

2 Epistemic Independence

At the outset we characterized NICs as conditional surface structures that relate clauses whose meanings are intuitively not conditionally related. I would like to try to pin down what it means for two propositions not to be conditionally related. Intuitively, some cases are clear: normally we would not expect the truth or falsity of propositions such as

$$\text{you are hungry } (P) \qquad \& \qquad \text{there are biscuits on the shelf } (Q)$$

to depend on one another. These two propositions should, in some sense, be independently true or false. But where? Not in the actual world where it is fixed whether P and whether Q. But perhaps rather in the mind of an agent: it is very plausible to assume that under normal circumstances a rational agent simply does not believe in any conditional connection between these propositions. A belief in a conditional connection between P and Q would be a belief that links the truth or falsity of one proposition to the truth or falsity of the other. Put the other way around, we might then say that P and Q are *epistemically independent* for an agent (in a given epistemic state) if learning one proposition to be true or false (where this was not decided before) is not enough evidence to decide whether the other proposition is true of false (where this was not decided before).

To make this idea more intelligible, gentle formalization helps. Take a set W of possible worlds, propositions $P, Q \subseteq W$ and an agent's epistemic state $\sigma \subseteq W$ of worlds held possible. We write $\overline{P}$ for $W \setminus P$, the negation of proposition P. We say that the agent holds P possible and write $\Diamond_\sigma P$ or, dropping the obvious index, $\Diamond P$ iff $\sigma \cap P \neq \emptyset$. Now, suppose we assume a strict conditional analysis and say that a conditional $P \rightarrow Q$ is held true in σ iff $\sigma \cap P \subseteq Q$. We could then say that an agent does not believe in any sort of conditional relationship between P and Q if he does not hold true any conditional $X \rightarrow Y$ for $X \in \{P, \overline{P}\}$ and $Y \in \{Q, \overline{Q}\}$. This would mean that an agent does not see a conditional relationship between P and Q iff

$$\forall X \in \{P, \overline{P}\} \ : \ \forall Y \in \{Q, \overline{Q}\} \ : \ \Diamond(X \cap Y). \tag{1}$$

This is indeed the standard notion of *logical independence* between propositions. For our purposes, however, this standard notion is too strong, for it commits the agent to believe in the possibility of P, $\overline{P}$, Q and $\overline{Q}$. This does not seem necessary: I may know that there are biscuits on the shelf, but still doubt any conditional relationship between that fact and your appetite. (We will soon see another formal reason why we should relax this criterion.)

So here is another attempt. We said that, for P and Q to be epistemically independent, learning that P is true or false should not decide whether Q is true or false. So, for $X \in \{P, \overline{P}\}$ and $Y \in \{Q, \overline{Q}\}$, if Y is held possible, then it should also be held possible after X is learned. That is to say:

$$\Diamond Y \rightarrow \Diamond(X \cap Y).$$

But again this only holds if X itself is at least possible, for otherwise, we again smuggle in a belief in the possibility of X via the definition of epistemic independence. We also aim for symmetry, so we take

$$(\Diamond X \wedge \Diamond Y) \rightarrow \Diamond(X \cap Y) \tag{2}$$

as our intuitive definition of epistemic independence of propositions P and Q: if P was not believed true or false before, then learning P is not enough to establish a belief in Q or $\overline{Q}$ where there was no such belief before.

2.1 Epistemic Independence and Conditionals

We noted as an explanandum that NICs seem to entail their consequent proposition. We are now in a position to account for this intuition. Let's boldly assume a material or strict implication analysis of conditionals even for NICs, but let's evaluate the conditionals on the epistemic state σ of a speaker. So, if a speaker says 'If P, Q', we may infer that, if he spoke truthfully, his epistemic state is such that $\sigma \cap P \subseteq Q$. But if we have reason to assume that at the same time the same speaker actually does not believe in a conditional relationship between P and Q, we may infer even more, namely that the speaker either believes in the falsity of P or the truth of Q. This is so, because if $\Diamond P$ and $\Diamond\overline{Q}$, then by epistemic independence we have $\Diamond(P \cap \overline{Q})$ which contradicts $\sigma \cap P \subseteq Q$. Consequently, if we furthermore have reason to assume that the speaker considers it at least possible that the antecedent proposition is true, as seems uncontroversial for all cases of NICs that I have seen so far, we may conclude that the speaker actually believes Q. Whence, I propose, the feeling of entailment: a speaker who (i) speaks truthfully in asserting 'If P, Q', (ii) considers P and Q epistemically independent and (iii) considers P at least possible *must* believe in Q.

2.2 Epistemic Independence and Probabilistic Independence

It is perhaps necessary to address an obvious worry. It might seem as if the notion of epistemic independence as defined in (2) appears out of thin air; an arbitrary formal tool shaped and designed to do exactly what we want it to. This is not so. In fact, it is the exact non-probabilistic counterpart of the notion of probabilistic independence of events in probability theory. In probability theory two events or propositions P and Q are said to be *probabilistically independent* given a probability distribution $\Pr(\cdot)$ iff

$$\Pr(P \cap Q) = \Pr(P) \times \Pr(Q). \tag{3}$$

If we identify the probability distribution $\Pr(\cdot)$ with the agent's subjective probabilistic beliefs over worlds W

$$\Pr : W \to [0;1], \quad \Pr(P) = \sum_{w \in P} \Pr(w), \quad \Pr(W) = 1$$

and therefore equate the epistemic state σ of the agent with the support of the probability distribution $\Pr(\cdot)$ as usual

$$\sigma = \{w \in W \mid \Pr(w) \neq 0\}$$

we can show that probabilistic independence of propositions P and Q as in (3) entails epistemic independence as in (2).[3]

[3] Proof: First, we establish that if $\Pr(P \cap Q) = \Pr(P) \times \Pr(Q)$, then for arbitrary $X \in \{P, \overline{P}\}$ and $Y \in \{Q, \overline{Q}\}$ it holds that $\Pr(X \cap Y) = \Pr(X) \times \Pr(Y)$. From the three arguments needed, it suffices

The converse, however, is not the case. Epistemic independence does not entail probabilistic independence. It may be the case that proposition P is not enough (evidence, support, information) to decide whether Q is true or false, but still learning that P is true, for instance, makes Q more or less likely. That is to say that the suggested notion of epistemic independence is indeed the exact non-probabilistic counterpart of probabilistic independence: we only care about believing and holding possible, not likelihoods of propositions.

This makes for a further argument why the stronger notion of logical independence of propositions in (1) may indeed be deemed too strong also on formal grounds. Instead of defining probabilistic independence of P and Q as in (3) we could, perhaps, have required that for X and Y as before

$$\Pr(X|Y) = \Pr(X) \quad \& \quad \Pr(Y|X) = \Pr(Y). \tag{4}$$

This is only defined if all X and Y have non-zero probability. This stronger notion entails, but is not entailed by the standard notion of probabilistic independence given in (1). It is obvious that logical independence as defined in (1) is the non-probabilistic counterpart of this stronger, non-standard definition of probabilistic independence in (4), in the same sense that epistemic independence (2) is the non-probabilistic counterpart of standard probabilistic independence (3).[4]

3 Information in Decision Contexts

It remains to be explained why the antecedents in NICs apparently serve to classify or conditionalize a speech-act whose felicity, in particular its relevance, is guaranteed only under the premiss that the antecedent is true. As stated earlier, I would like to defend that we do not have to assume implicit performatives or abstract speech-act operators in the semantics to account for this intuition. I can only sketch my arguments for this thesis in rough outline here. The brief indications I will give apply in an obvious manner to relevance conditionals like (1) and (3), but do not extend directly to meta-communicative (2a) or discourse-structuring conditionals (2b). Still, I believe that a similar story can be told for these too.

On a certain level of abstraction, we may think that most, if not all, reasonable communication takes place against the background of a decision problem of the hearer. A *decision problem* is a quadruple $\delta = \langle W, u, \Pr, Act\rangle$ where W is a set of relevant situations or world states, u is a utility function from W to reals, $\Pr(\cdot)$ is a probability distribution over W and Act is a partition of W (a finite set of propositions) representing the hearer's possible future actions. We can be very liberal interpreting what actions are. Actions may be very abstract *epistemic actions*: interpretations, here conceived as the (deliberate) adoption of

to give just one, as the others are similar. So assume that $\Pr(P \cap Q) = \Pr(P) \times \Pr(Q)$ and derive that $\Pr(P \cap \overline{Q}) = \Pr(P) \times \Pr(\overline{Q})$: $\Pr(P \cap \overline{Q}) = \Pr(P) - \Pr(P \cap Q) = \Pr(P) - (\Pr(P) \times \Pr(Q)) = \Pr(P) \times (1 - \Pr(Q)) = \Pr(P) \times \Pr(\overline{Q})$. Next, assume that $\Pr(X \cap Y) = \Pr(X) \times \Pr(Y)$ and that $\Diamond P$ and $\Diamond Q$. That means that $\Pr(X), \Pr(Y) > 0$. Hence, $\Pr(X \cap Y) > 0$, which is just to say that $\Diamond(X \cap Y)$. □

[4] A final remark in defense of the proposed notion of epistemic independence: van Rooij (2007) noticed that epistemic independence is equivalent to Lewis's (1988) notion of orthogonality of questions whether P and whether Q. In the same paper, van Rooij uses epistemic independence/orthogonality to account for the strengthening of conditional presuppositions to unconditional ones. The underlying idea is the same as the one presented here.

beliefs about what was meant with an utterance. Thus conceived we may think of a background decision problem as an abstraction over the notion of a question under discussion: reasonable talk exchange addresses an issue "what shall I do?" of which a special case is the question "what shall I believe?".

In a decision problem δ, a rational agent is predicted to choose an action $A \in Act$ that maximizes expected utility which is defined as

$$\mathrm{EU}(A) = \sum_{w \in A} \Pr(w) \times u(w).$$

Information that proposition P is true is processed via Bayesian update to yield an updated decision with $\Pr(\cdot|P)$ instead of $\Pr(\cdot)$. The expected utility of an action A after learning that P

$$\mathrm{EU}(A|P) = \sum_{w \in A} \Pr(w|P) \times u(w).$$

may differ from $\mathrm{EU}(A)$, of course. Also, the actions with the highest expected utility before learning and after learning need not be identical. This is neither new nor surprising. It simply means that in a concrete decision problem mere information flow may have a particular impact on the decision of the hearer, i.e. relate to the question under discussion.

The interesting question then is, when and how is P *relevant* to the decision problem δ? In other words, when does P count as an answer to the question under discussion? I suggest a weak notion of answerhood: P is relevant to decision problem δ iff it is a non-trivial argument for some action in δ. For formalization of the notion of an argument, define the *change of expected utility* of action A given P as

$$\mathrm{CEU}(A, P) = \mathrm{EU}(A|P) - \mathrm{EU}(A)$$

and say that P is an argument for A iff

(i) $\forall B \in Act\ \mathrm{CEU}(A, P) \geq \mathrm{CEU}(B, P)$ and

(ii) $\exists B \in Act\ \mathrm{CEU}(A, P) > \mathrm{CEU}(B, P)$.

Equivalently, we could say that P is relevant iff there are actions A, B in Act such that $\mathrm{CEU}(A, P) \neq \mathrm{CEU}(B, P)$. Yet the lengthier formulation makes clear in which sense relevant information is an answer to the question under discussion: if P is relevant it sets apart at least one action.

In real life we would certainly not want to assume that a speaker at all times knows exactly what drives a hearer, what he is concerned with and what he wants. In other words, the background decision problem of the hearer may at times not be known by the speaker, either in its entirety or in some relevant detail. Consequently, it may not always be the case that a piece of information is relevant to all the decision problems that the speaker might consider possible issues for the hearer. Similarly, a hearer may not be sure in the context of which speaker-conceivable decision problem he should interpret statements by the speaker. Enter relevance conditionals like (1) or (3).

If I'm in a shop, the information that the name of the woman helping me find the dressing room is Jill may be relevant in a million ways. (Does she want me to ask for her phone number?) If you pay me a visit and I just say out of the blue that there are biscuits on the

shelf, your most natural reaction is perhaps a stunned: "So what?". Yet, with the NICs in (1) and (3) the case is entirely different. We now know much better in which way we ought to process the information in the consequent, in which sense it is relevant.

How can this intuition be made a little more precise? Suppose that the hearer who interprets the NIC "If P, Q" adds to his stock of knowledge, Ramsey-style, the proposition P. He realizes that the information that Q is true is not restricted to P-worlds, because P and Q are epistemically independent etc. But he also realizes that the speaker may believe that all the P-worlds are worlds in which the hearer has a particular decision problem. That means that the hearer, after adopting P for the evaluation of Q, is in a particular context in which, plausibly, the information Q is indeed relevant: if you are hungry, the information that there are biscuits on the shelf simply increases the expected utility of one action (going over and having some biscuits) more than some other action (ordering a pizza) and therefore is an argument for that action, hence relevant for that decision problem, no matter what exactly the beliefs and utilities of the hearer are.

Though sketchy, this outline should make clear that we can dispense with conditional assertions in the case of relevance conditionals. Instead of saying that, for reasons of unclear relevance, the consequent proposition is asserted if (and, presumably, only if) the antecedent is true, we say that the antecedent (dynamically, if you wish) takes us to a context in which the consequent proposition is relevant. Being relevant then means: relating or arguing to a point. The twist is that relevance does not apply to assertions, but to information.

So what if the antecedent is false? In that case we predict that still a normal conditional has been asserted from which the hearer learns via independence that the consequent proposition is true, but this information may not be relevant, or it is relevant in a different (kind of) decision problem where it relates to and argues for a different point. Hence, by indicating which context the consequent information is to be understood in the speaker expresses that this is the effect (think: point, argument) of the information in the consequent; if the necessary context is not actual, this effect simply doesn't come about. What sounds like a conditional speech-act, is just information processed in the context of a decision problem.

References

Johan van der Auwera. Conditionals and speech acts. In Elizabeth Closs Traugott, Alice ter Meulen, Judith Schnitzer Reilly, and Charles A. Ferguson, editors, *On Conditionals*, pages 197–214. Cambridge University Press, 1986.

Keith DeRose and Richard E. Grandy. Conditional assertions and 'biscuit' conditionals. *Noûs*, 33(3):405–420, 1999.

Susanne Günthner. Wenn-Sätze im Vor-Vorfeld: Ihre Formen und Funktionen in der gesprochenen Sprache. *Deutsche Sprache*, 3:209–235, 1999.

David Lewis. Relevant implication. *Theoria*, 54:162–174, 1988.

Robert van Rooij. Strengthening conditional presuppositions. To appear in Journal of Semantics, 2007.

Muffy E. A. Siegel. Biscuit conditionals: Quantification over potential literal acts. *Linguistics and Philosophy*, 29:167–203, 2006.