

Rethinking Gibbard's Riverboat Argument

Karolina Krzyzanowska
Faculty of Philosophy, University of Groningen
k.h.krzyzanowska@rug.nl

Sylvia Wenmackers
Faculty of Philosophy, University of Groningen
s.wenmackers@rug.nl

Igor Douven
Faculty of Philosophy, University of Groningen
i.e.j.douven@rug.nl

Abstract

According to the Principle of Conditional Non-Contradiction (CNC), conditionals of the form “If p , q ” and “If p , not q ” cannot both be true, unless p is inconsistent. This principle is widely regarded as an adequacy constraint on any semantics that attributes truth conditions to conditionals. Gibbard has presented an example of a pair of conditionals that, in the context he describes, appear to violate CNC. He concluded from this that conditionals lack truth conditions. We argue that this conclusion is rash by proposing a new diagnosis of what is going on in Gibbard's argument. We also provide empirical evidence in support of our proposal.

1. Gibbard's riverboat argument. According to the principle of Conditional Non-Contradiction (CNC), conditionals with the same antecedent and contradictory consequents cannot both be true, unless the antecedent is inconsistent.^{1,2} We follow the mainstream in regarding this principle as an adequacy constraint on any truth-conditional semantics for conditionals. That immediately rules out the once-popular material conditional account, which assigns to a conditional the truth conditions of its material counterpart; given this semantics, “If p , q ” and “If p , not q ” are both true if p is false. But surely CNC leaves many other truth-conditional semantics in the running? An old but still influential argument by Gibbard [1981] purports to show that

¹This paper is concerned exclusively with indicative conditionals. We use “conditional” to refer to indicative conditionals throughout.

²Some authors use the label CNC to refer to what we call CNC minus the proviso that the antecedent be consistent. The principle as stated here is then referred to as “Restricted Conditional Non-Contradiction.” See for instance Unterhuber [2013, Ch. 3], which also contains a thorough discussion of the status of both principles in the context of various semantics for conditionals.

the answer is negative. In Gibbard's view, any truth-conditional semantics that validates CNC is bound to get things wrong—from which he concludes that conditionals do not have truth conditions.³

Central to Gibbard's argument is this story:⁴

Sly Pete and Mr. Stone are playing poker on a Mississippi riverboat. It is now up to Pete to call or fold. My henchman Zack sees Stone's hand, which is quite good, and signals its content to Pete. My henchman Jack sees both hands, and sees that Pete's hand is rather low, so that Stone's is the winning hand. At this point, the room is cleared. A few minutes later, Zack slips me a note which says "If Pete called, he won," and Jack slips me a note which says "If Pete called, he lost." I know that these notes both come from my trusted henchmen, but do not know which of them sent which note. I conclude that Pete folded. (Gibbard [1981:231])

Gibbard argues that if

(1) If Pete called, he won.

and

(2) If Pete called, he lost.

have a truth value at all, they must both be true. This is because Zack and Jack both appear warranted in their assertions, and their warrants do not rest on any false beliefs about relevant matters of fact. But CNC rules that (1) and (2) cannot both be true. Truth-conditional semantics for conditionals that validate CNC—as we granted they must—cannot get this right.

Or rather, Gibbard does see an escape route for the advocates of such semantics, namely, to claim that conditionals have their meanings relative to speakers' epistemic states, so that (1) does not express the same proposition when it is uttered by Zack as when it is uttered by Jack. Supposing this to be true, we would equivocate by taking CNC to apply to (1) and (2). But—Gibbard thinks—this route leads to a kind of context-sensitivity that is hard to uphold. After all, we do seem able to make sense of (1) and (2) even if we are ignorant of the epistemic states that the authors of these notes were in at the time of their writing.

We already granted CNC. We are also willing to grant Gibbard that if one of (1) and (2) is true, then so is the other.⁵ If we still do not accept the conclusion that conditionals must lack truth conditions, then this is because we disagree with Gibbard's

³The same position is defended on independent grounds by, among others, Adams [1975], Appiah [1985], and Edgington [1995].

⁴For readers who are—like us—unfamiliar with the game of poker, we note that the only important things to understand are that if Pete *calls* he decides to continue to play but thereby risks losing money (though also has a chance of winning), and that if he *folds* he quits and does not risk losing any money (but also cannot win).

⁵We are granting this even though we are not convinced by Gibbard's argument for it. True, neither Jack's nor Zack's assertion rests on a mistaken belief, but Jack, in contrast to Zack, still lacks knowledge of a seemingly highly relevant matter of fact, to wit, what cards Pete is holding. While Zack may not have

appraisal of the escape route he mentions. Why should it be more objectionable to claim that a conditional can express different propositions in the mouths of different speakers than to claim that, for instance,

(3) I am the owner of this house.

expresses different propositions in the mouths of different speakers, and may be true when said by Sally and false when said by Paul? According to Gibbard, the cases are not parallel because we can rely on fairly straightforward rules for detecting the relevant contextual information necessary to interpret sentences like (3)—“I refers to the speaker, ‘this’ to whatever the speaker points at”—but there do not seem to exist similarly straightforward rules for interpreting putatively context-sensitive conditionals.

We dispute Gibbard’s assessment of the situation and believe that, pace Gibbard, interpreting context-sensitive conditionals is not so different from interpreting sentences containing indexicals. This is not to deny that we can make sense of (1) and (2) even if we have no idea which note is from Zack and which from Jack, nor know much about Zack’s and Jack’s belief states. But in the same way we can make sense of (3) if that is written on a note of unknown origin, and we do not know the referents of “I” or “this house.” What is true is that in such a case we cannot *fully* interpret (3), but we still get *something* from it; we get from it that whoever wrote the message claims to be the owner of whichever house “this house” is meant to refer to. In some contexts, we might even be able to say with some confidence who the author is and to which house he or she is referring. In our view, much the same is true for (1) and (2). Having been given the notes with (1) and (2) on them, but not knowing who gave us which note and what the persons’ epistemic states were, we can still get *something* out of those messages. To be sure, one gets *more* out of the messages if one does know who handed one which note, and what the relevant beliefs of the authors were at the time of writing. But that is as with the note with (3) on it: one gets more out of that if one does know the intended referents of “I” and “this house” than if one does not.⁶

To back up this view, we want to present a truth-conditional semantics for conditionals that renders both (1) and (2) true, given the specifics of the above story. We are by no means sure that the semantics is correct, nor, however, do we at this point see any shortcomings. At a minimum, the position should be on the map, if only because it allows us to agree with Gibbard that both (1) and (2) are true while keeping CNC on board *and* avoiding the untoward consequences that Gibbard envisions for any semantics that factors a speaker’s epistemic state into the truth conditions of conditionals. Moreover, we report some first experimental evidence favoring our proposal.

any *mistaken* beliefs about these cards, his ignorance might still lead him to assert something false—even warrantably so, given the quality of the evidence that *is* in his possession. See, in the same vein, Lycan [2001:169].

⁶See Krzyżanowska [2012] for more on this.

2. Toward a new semantics for conditionals. The semantics that we want to propose takes its cue from an old idea according to which a conditional is true iff its consequent follows from its antecedent. This idea was meant to do justice to the broadly felt intuition that there must be some kind of internal connection between a conditional's antecedent and its consequent if that conditional is to count as true. If this idea, which dates back to ancient logicians and has been defended by many philosophers since (see Sanford [1989, Chs. 1 and 2]), enjoys little popularity today, then that is because its advocates have wanted to understand the notion of a consequent's following from an antecedent strictly in terms of (classical) deductive inference.⁷ It is not hard to come up with conditionals that pre-theoretically appear true yet whose consequent does not follow logically from their antecedent. Our proposal aims to avoid this problem by doing justice to the fact that a consequent's following from an antecedent may be understood in terms of a number of different inferential connections, including but not limited to deductive inference.

In doing so, we are following the lead of Douven and Verbrugge [2010], who propose a typology of conditionals based on the type of inference embodied by a conditional. They first distinguish between certain or deductive inferential (DI) conditionals and uncertain conditionals, and then go on to subdivide the latter into abductive inferential (AI) conditionals and inductive inferential (II) conditionals. As the names suggest, the distinction between these types of conditional is made on the basis of the type of inferential connection that holds between a conditional's antecedent and its consequent, where the connection may be deductive, abductive, or inductive.⁸ The deductive consequence relation is familiar from standard logic courses, and the non-deductive consequence relations have also received a fair amount of attention in the literature; see, for instance, Kyburg and Teng [2001] and Gabbay and Woods [2005]. Still, there is much less agreement on how abductive and inductive inference are best explicated than agreement on how deductive inference is to be explicated. For present purposes, we need not commit ourselves to any specific understanding of abductive or inductive inference; it will suffice to say that q is an abductive consequence of p (given the background premises) iff q best explains p (in light of the background premises), and that q is an inductive consequence of p (given the background premises) iff q follows with high statistical probability from p (in light of the background premises).

⁷Or rather, it enjoys little popularity among contemporary philosophers. At least in some of their writings, psychologists working in the mental models tradition seem to still be endorsing this position; see, for instance, Johnson-Laird and Byrne [2002], who say that for a conditional to be true “the consequent *has to occur given the antecedent*” (p. 649, italics ours).

⁸As the explicit reference to *inferential* conditionals also suggests, Douven and Verbrugge see themselves as being concerned with one type of conditionals only, to wit, inferential conditionals. Linguists tend to distinguish these from what they call “content conditionals.” Douven and Verbrugge [2010:303] remain neutral on the question of whether what linguists identify as content conditionals constitute a genuinely different class of conditionals. Here, we do so as well, and more generally remain neutral on how general our proposal is. Even if not all conditionals are inferential conditionals, the ones figuring in Gibbard's story clearly are, so for our solution to the puzzle that story raises it is immaterial whether our semantics is fully general or pertains only to one important class of conditionals.

Douven and Verbrugge also distinguish between contextual and non-contextual versions of these types of conditionals: for contextual versions, the inferential relation between antecedent and consequent rests partly on background premises that are assumed to hold in the context in which the conditional is asserted or evaluated. For example, considered on its own, q may fail to explain p , but in light of all that one believes, q may be the best explanation of p . If so, then “If p , q ” qualifies as a contextual AI conditional, where the relevant context is provided by one’s current belief state.

What Douven and Verbrugge do not note in their 2010 paper (and what is also not important for their purposes) is that there can be an inferential connection between antecedent and consequent which involves inferences of more than one of the aforementioned types. For example, in some conditionals the consequent may follow from the antecedent via an abductive step *and* a deductive step. In cases of this kind, we may say that the consequent is a mixed consequence of the antecedent. It is important to note that as soon as either abductive or inductive steps are involved, the conditional is to be grouped with the uncertain inferential conditionals.

While, to our minds, the above typology of consequence relations is directly relevant to the semantics for conditionals, it may be difficult to capture this relevance by the standard model-theoretic means, if only because there is currently no satisfactory model-theoretic characterization of best explanation. But there is an alternative tradition in semantics, one that uses proof theory instead of model theory for the purposes of explicating meaning. This tradition is best known for its claim that the meanings of the logical constants are given by the standard introduction and elimination rules for these constants.

Our suggestion is *not* that the standard introduction and elimination rules for the conditional operator from propositional logic give us the semantics for the conditional; that would amount to endorsing the material conditional account, which we have already dismissed. Rather, the idea is to state truth conditions for conditionals directly in terms of the various mentioned consequence relations. Doing so helps to amend in two ways the traditional idea that was our starting point, to wit, by relaxing the requirement that a conditional’s consequent follows *deductively* from its antecedent for the conditional to be true, and by acknowledging the role the background plays in determining whether or not a conditional’s consequent follows from its antecedent. To be precise, the proposed semantics is that

Definition 1 *A speaker S ’s utterance of “If p , q ” is true iff (i) q is a consequence—be it deductive, abductive, inductive, or mixed—of p in conjunction with S ’s background knowledge, (ii) q is not a consequence—whether deductive, abductive, inductive, or mixed—of S ’s background knowledge alone but not of p on its own, and (iii) p is deductively consistent with S ’s background knowledge or q is a consequence (in the broad sense) of p alone.*

Note that we are requiring background *knowledge* and not merely background *beliefs*: it would be counterintuitive to designate the utterance of a conditional as true if its

consequent followed (in any of the mentioned senses) from its antecedent in conjunction with false beliefs that the speaker may have.

Clauses (ii) and (iii) are meant to ensure that the antecedent is not redundant in the derivation of the consequent, respectively, that the consequent does not follow trivially from the antecedent plus background knowledge. Without them, one could still have true conditionals without any intuitive inferential connection between antecedent and consequent. As for clause (ii), note that the proposal would be too restrictive if we demanded simply that the consequent not follow from the background knowledge alone. While we do want to keep from qualifying as true

- (4) If Hillary Clinton runs for president in 2016, the earth weighs more than 2 kilograms.

we do *not* want to keep from qualifying as true

- (5) If the earth weighs more than 3 kilograms, then it weighs more than 2 kilograms.

The consequent of these conditionals follows (in almost any context) from background knowledge alone. However, in (5) it also follows from the antecedent, which is enough to ensure the intuitively required inferential link, whereas in (4) antecedent and consequent have nothing to do with each other, which at least on our hypothesis accounts for the felt falsity of this sentence. The proposal would also be too restrictive if clause (iii) simply required that the antecedent be deductively consistent with the background knowledge. The conditional

- (6) If the UK is ruled by a king, it is a monarchy.

seems true, even if its antecedent is inconsistent with background knowledge. By contrast,

- (7) If the UK is ruled by a king, Hillary Clinton will run for president in 2016.

seems false, notwithstanding the fact that its consequent follows deductively from its antecedent together with the background knowledge that the UK is (currently) ruled by a queen; whether Hillary Clinton will run for president in 2016, has, as far as we can see, *nothing* to do with whether the UK is ruled by a king.

As the proposal stands, conditionals with a necessarily true consequent as well as those with a necessarily false antecedent still qualify as true. Thus, clauses (ii) and (iii) do not quite ensure what they are intended to ensure. But perhaps that is as it should be. We do not have strong intuitions about whether, say,

- (8) If $2 + 2 = 5$, Hillary Clinton will run for president in 2016.

is true or false. Those who do have a strong intuition that it is false could still consider dealing with this type of conditional by appealing to Gricean pragmatics.⁹ More

⁹Some might think that Gricean pragmatics on its own is capable of accounting for the intuition that there should be a link between a conditional's antecedent and consequent if that conditional is to count as true. However, see Douven [2008] for an argument to the effect that no currently known pragmatic principle can do the required job.

generally, we would not be hugely surprised if clauses (ii) and (iii) needed further fine-tuning.

For the purposes of shedding new light on Gibbard's riverboat argument, the important clause of Definition 1 is the first one, which may be considered to be the core of the current proposal. This clause raises some tangled issues by itself, in particular in relation to II conditionals, that is, conditionals whose consequent is an inductive consequence of their antecedent. To mention a pressing one, few might want to say that a speaker informed about the circumstances of a fair and large lottery could truly assert the II conditional "If you buy a ticket, you will lose," however many tickets there are in the lottery and however low the chances are of winning the lottery (Lowe [1996]). Of course, this may just mean that some clause is to be added to our proposal. Alternatively, one might try to argue that the intuition that "If you buy a ticket, you will lose" cannot be truly asserted is to be explained in terms of this conditional's not being assertable in the first place, which in turn may have an explanation along Gricean lines (Douven [2012]). Given that no II conditionals are involved in Gibbard's story, we flag this potential difficulty here only to set it aside for future research.

Before applying our proposal to Gibbard's argument, we would like to point out some attractive features of Definition 1 that are independent of that argument. First, note that it has no difficulty blocking the paradoxes of material implication. Because q may be true without there being any inferential connection between it and p , "If p, q " need not be true if q is true. Nor does the truth of the conditional follow from the mere falsity of p , as that does not either ensure the existence of an inferential connection between p and q .

Furthermore, the semantics gets the Or-to-if principle right precisely in the kind of cases in which it is intuitively right. According to this principle, "If p, q " can be inferred from "Not- p or q ." The principle is generally thought to be compelling, but it is also known that, as it stands, the principle is validated only by the material conditional account.¹⁰ Now, if q follows from the background knowledge alone in a context, then so does "Not- p or q ."¹¹ However, given clause (ii) of Definition 1, it still does not follow that "If p, q " is true in that context; it is *not*, unless q also follows from p alone. Similarly if not- p follows from the background knowledge. Then even though "Not- p or q " follows from that background knowledge as well, it does, given clause (iii), *not* follow that "If p, q " is true in the given context; it is *not*, unless q follows from p alone. Note, however, that in these two kinds of situations it is not intuitively all right to infer "If p, q " from "Not- p or q ": if we know that the butler did it, we do not want to infer that if the butler did not do it, then the maid did; nor do we want to infer this if we know that the maid did it. On the other hand, in the kind of case in which neither not- p nor q follows from the background knowledge though their disjunction does, it is perfectly all right to apply the Or-to-if principle. And given Definition 1, "If p, q " is

¹⁰See Stalnaker [1975].

¹¹We are assuming knowledge to be closed under classical entailment.

true in that case, for given “Not- p or q ,” there is an inferential connection—to wit, via Disjunctive Syllogism—from not- p to q .

Finally, the new semantics validates the so-called Import–Export principle, according to which “If p and q , r ” and “If p , then if q , r ” are equivalent. To see this, let “If p and q , r ” be given. Then r can be inferred (in the broad, generalized sense) from the conjunction of p and q ; let Δ be the name of one way to derive r from the conjunction of p and q . Then from p we can infer that r can be inferred from q . After all, given p as a premise, we can assume q and form the conjunction of p and q . Then we can use Δ to derive r from that conjunction. Discharging q yields that “If q , r ” can be inferred from p .¹² And this means that “If p , then if q , r ” is true. Conversely, let “If p , then if q , r ” be given. Then from p we can infer that r can be inferred from q . Supposing the conjunction of p and q , we first infer from p that r can be inferred from q and then use that and q to infer r . Thus we infer r from the conjunction of p and q , which is enough for the truth of “If p and q , r .”

One may wonder, however, what remains on the present proposal of the intuition that modus ponens is a valid rule of inference for conditionals (and not just for the material conditional). By relaxing the requirement of previous semantics that the inferential connection between antecedent and consequent be deductive, our semantics makes room for the possibility that a true conditional has a true antecedent and a false consequent. After all, what distinguishes deduction from induction and abduction is that deduction guarantees the truth of any conclusion reached on the basis of true premises. But this means that modus ponens may lead from true premises to a false conclusion, and hence that the rule is not valid in the classical sense of guaranteeing preservation of truth.

In response, first note that the validity of modus ponens has been challenged on independent grounds by McGee [1985] and Lycan [2001, Ch. 5]. More importantly, even if abduction and induction are not *guaranteed* to preserve truth, we rely on them in daily practice because we take them to be reliable guides to the truth; that is, we trust them to preserve truth with high probability. If we are right to trust these modes of inference, then, supposing our semantics, in the vast majority of cases in which we apply modus ponens to a conditional with a true antecedent, its consequent will be true as well. That may be all there is to the intuition that modus ponens is a classically valid rule of inference—and it may be all that matters for practical purposes.

To return to Gibbard’s story now, and see what follows for it from Definition 1, first consider a variant of that story. Sly Pete and Mr. Stone are playing poker again, with Zack and Jack again doing their business. Once more, it is up to Pete to call or fold, who also in this story has been signalled the content of Mr. Stone’s hand by Zack. In contrast to the original story, however, Pete now has the upper hand. Zack does not know this, but Jack, who has again seen both players’ hands, does. After the

¹²This assumes the validity of Conditional Proof for the conditional, that is, that we may conclude “If p , q ” if we can infer (still in the relevant broad sense) q from p . But Conditional Proof is trivially valid on the present account, which explicitly makes it sufficient for the truth of a conditional that its consequent can be inferred from its antecedent.

room has been cleared, both Zack and Jack sincerely assert (1). Although by doing so, they both assert something true, what makes Zack's assertion true is not the same as what makes Jack's assertion true. Relative to Zack's background knowledge, the best explanation for Pete's calling—supposing he called—is that he held cards that were better than Mr. Stone's, which Pete knew thanks to Zack's fraudulent act. From this in conjunction with the antecedent of (1), it follows deductively that Pete won, so that, from Zack's perspective, the consequent of (1) is a mixed consequence of the antecedent of this conditional. That is what makes Zack's assertion of (1) true. On the other hand, relative to Jack's background knowledge, that Pete won follows logically from his calling together with the information about both players' hands and the rules of poker. *That* is what makes Jack's assertion of (1) true.

The proposed diagnosis of what is going on in Gibbard's original story will now be unsurprising. From Jack's perspective, "Pete lost" obviously follows deductively from "Pete called," which is what makes his assertion of (2) true. This does not mean that Zack's utterance of (1) must be false. To the contrary, as in the above story, from Zack's perspective Pete's calling is best explained by his having a winning hand, from which, in conjunction with the antecedent, it follows deductively that Pete won—and this is what makes Zack's assertion of (1) true. That is to say, according to Definition 1, both Jack and Zack assert something true.^{13,14}

In summarizing Gibbard's argument in Section 1, we said that, by taking the escape route suggested by Gibbard, his story would no longer present a violation of CNC. We took that escape route, but it could be objected that, strictly speaking, we still have a violation of CNC. For, as it stands, CNC assumes that one can determine merely on the basis of form whether a conditional in the mouth of one speaker is inconsistent with a conditional in the mouth of another speaker. And the above proposal has not done anything to alter the form of either (1) or (2). As will be clear, we take this to overlook the crucial role that the background knowledge of a speaker plays in determining the interpretation of a conditional. To make this role explicit, it suffices to reformulate CNC as the principle that utterances of instances of "If p , q " and "If p , not q " cannot both be true in conjunction with the same body of knowledge. Put this way, the conditionals in Gibbard's story no longer constitute a violation of CNC.

¹³A number of authors (e.g., Bennett [2003:85]) have presented "symmetric" versions of Gibbard's story which involve pairs of conditionals that are both of the same inferential type. But given that in those versions the personae asserting the conditionals still have different background knowledge, there is no problem acknowledging that the conditionals are both true.

¹⁴It might be said that accounts like ours, that relativize the truth of a conditional to the speaker's background knowledge, fail to address the intuition that Zack and Jack disagree with each other. In our view, however, their disagreement is only apparent. If Jack and Zack were to talk to each other, they would quickly realize that they have different information, which might lead them to share their information. Probably, that would make Zack change his mind. But this would not really amount to retracting his earlier assertion. Rather, we would expect Zack to say something like the following once Jack had revealed his information: "What I meant to say is that if Pete called, that would have indicated that he had the upper hand. But now that Jack told me that in fact Mr. Stone had the upper hand, I'm sure Pete did not call."

To come to what is arguably the main challenge of Gibbard's argument, we note that while our view makes the interpretation of conditionals relative to speakers' background knowledge, it does not lead to a relativism that must leave one completely clueless as to the interpretation of any given conditional. Not knowing who wrote which note, we still understand that Pete's losing is a consequence of his calling relative to one of the henchmen's background knowledge, and Pete's winning is a consequence of his calling relative to the other henchman's background knowledge. Naturally, we understand more if we know who wrote which note, and if we know something about his background knowledge. For instance, knowing that (1) comes from Zack, and that Zack signalled Mr. Stone's hand to Pete but did not see Pete's hand, it should not take too much effort to figure out that Pete's winning is not just *a* consequence of his calling, but more specifically a *mixed* consequence, involving an abductive inferential step. Still, Gibbard or anyone else defending his line would be hard pressed to argue how any of this is essentially different from reading a note with, for instance, (3) on it in a context in which one is ignorant of who wrote the note as compared to reading that note in a context in which one does know who wrote it. They would be equally hard pressed to argue that, on our account, one gets less from reading (1) or (2) than one does from a pre-theoretic viewpoint.

The main general point to be emphasized in this connection is that interpretation is not a matter of all or nothing. It is not as though we always either do or do not get what a speaker says; there are many shades of partial understanding in-between these extremes. By reading his henchmen's notes, the first-person narrator in Gibbard's story gets a lot, just not everything.

In short, the proposed semantics has several attractive features. There are important details to be filled in and it remains to be seen whether the proposal is able to pass other benchmark tests for semantics for conditionals. Most pressingly perhaps, can it explain the accumulated data on people's judgments of the probabilities of conditionals, or on how people reason with conditionals? And what does our proposal imply for the contested issue of whether conditionals express propositions? Concomitantly, what does it imply for Lewis' [1976] celebrated triviality arguments? We leave answering these questions for another occasion, and turn to something more modest, to wit, showing that at least our analysis of Gibbard's story as involving conditionals embodying different types of inferential connection between antecedent and consequent enjoys empirical support.

3. Inferential conditionals and evidential markers. Our diagnosis of what is going on in Gibbard's argument is so far mainly supported by our own pre-theoretic responses to (1) and (2) in the context provided by Gibbard's story. While in our opinion intuition-based approaches still have their value in philosophy, we have considerable sympathy for the recent trend to try and provide more robust, less subjective support for one's philosophical claims and analyses. This section presents an experiment that

is meant to offer precisely that kind of additional support for our analysis of Gibbard’s argument.

The experiment makes use of previous experimental work on so-called evidential markers in relation to conditionals (Krzyżanowska, Wenmackers, and Douven [2013]). In that work, we found that “probably,” when occurring in the consequent of a conditional, is a good marker of conditionals that are uncertain in Douven and Verbrugge’s [2010] sense. Specifically, inserting “probably” in the consequent of a conditional has a tendency to *raise* that conditional’s perceived assertability if the conditional is either an II or an AI conditional as well as a tendency to *lower* the conditional’s perceived assertability if the conditional is a DI conditional.

This finding suggests a straightforward empirical test of our diagnosis of Gibbard’s riverboat argument. After all, according to that diagnosis, “If Pete called, he lost” is a certain conditional relative to the background information available to Jack, and so, in view of the above finding, the insertion of “probably” in the consequent should make the conditional less fitting into a context in which that information is assumed. Conversely, according to the same diagnosis, “If Pete called, he won” is an uncertain conditional relative to the background information available to Zack, and so the insertion of “probably” in the consequent should make the conditional more fitting into a context in which *that* information is assumed.

3.1. Method

PARTICIPANTS

Two hundred and seventy-eight persons participated in the experiment. They were recruited via the CrowdFlower interface (<http://www.crowdflower.com>), which directed them to the Qualtrics platform (<http://www.qualtrics.com>) on which the experiment was run. The participants were paid a small amount of money in return for their cooperation. All participants were from Australia, Canada, New Zealand, the United Kingdom, and the United States. We excluded from the analysis 81 participants who failed two comprehension questions, 11 further participants who had indicated that they were non-native speakers of English, and 3 still further participants who had spent more than 30 minutes on the survey. This left us with 183 participants. The mean age of these participants was 35 (± 13); 115 participants (63 %) were women; 79 % indicated college or higher as their education level, 20 % indicated high school, and 1 % indicated a lower education level.

DESIGN

Every participant was given five different stories and was asked which of two conditionals fitted better in the context of the story. There were two versions of each story, a “certain” and an “uncertain” one. Each participant received only one version of a story. All the materials were presented on screen, with each story together with the associated question appearing on a different screen. The order of the story–question pairs was randomized and for each separate story it was randomly determined whether the agent received the certain or the uncertain version of that story.

MATERIALS AND PROCEDURE

All materials were in English, the participants' native language. One group of participants ($N = 83$) was offered the part of Gibbard's story containing the information available to Jack, who sees both players' hands (which creates a "certain" context), and another group ($N = 100$) was offered the part of Gibbard's story containing the information available to Zack, who sees only Mr. Stone's hand and signals its contents to Pete (which creates an "uncertain" context). Participants in the first group were then asked which of "If Pete called, he lost" and "If Pete called, he probably lost" they thought fitted best in the context. Participants in the second group were asked which of "If Pete called, he won" and "If Pete called, he probably won" they thought fitted best into the context. Both groups were given the options to answer that the first conditional fits best, that the second fits best, that both fit equally well, and that neither fits well.

More in detail, the first group was presented the following on screen:

Pete and Mr. Stone were playing poker yesterday. At some point, it was up to Pete to call (i.e., to continue playing, with the risk of losing if his cards are worse than Mr. Stone's) or to fold (i.e., to quit playing). Jack, who was in the same room, saw that, at that point, Mr. Stone's cards were better than Pete's.

Consider these sentences:

1. If Pete called, he lost.
2. If Pete called, he probably lost.

Of these sentences

- the first fits best into the context.
- the second fits best into the context.
- both fit equally well.
- neither fits well.

The corresponding story plus question presented to the second group was this:

Pete and Mr. Stone were playing poker yesterday. At some point, it was up to Pete to call (i.e., to continue playing, with the risk of losing if his cards are worse than Mr. Stone's) or to fold (i.e., to quit playing). Pete is a good poker player who does not like to lose. Moreover, he was cheating: his friend, Jack, was looking at Mr. Stone's hand and, using special signs, informed Pete about what cards Mr. Stone had.

Consider these sentences:

1. If Pete called, he won.
2. If Pete called, he probably won.

Of these sentences

- the first fits best into the context.
- the second fits best into the context.
- both fit equally well.
- neither fits well.

Because our own first impression of Gibbard's riverboat story was that it was somewhat difficult to comprehend, at least for people who are unfamiliar with the game of poker, we also included in our experiment DeRose's [2010] version of Gibbard's story, which is about a simpler game than poker but otherwise identical to Gibbard's story.¹⁵ As we did in the case of Gibbard's story, we split the information presented in DeRose's story into a part containing the information available to one player's helper—we called the player "Jim" in the survey question—who sees only what is in the hand of Jim's opponent—called "Steve" in the survey question—and signals its content to Jim, and a part containing the information available to Steve's helper who sees both players' hands. Again we asked one group ($N = 87$) to assess the fittingness in the former context of "If Jim plays, he will win" with and without "probably" in its consequent, and another group ($N = 96$) to do the same for "If Jim plays, he will lose" in the latter context. (See the appendix for the two versions of DeRose's story that we used.)

As an additional test of the correctness of our analysis of Gibbard's story and our hypothesis concerning what it shows about conditionals, we invented a number of stories that, supposing our analysis to be correct, closely parallel Gibbard's story in all relevant respects and that should therefore yield similar results with regard to the effect of inserting "probably" in the consequent of the conditionals that we supposed to parallel the conditionals in Gibbard's story. Specifically, we came up with three stories, one about a certain Mr. Smith, one about a poisoned apple, and one about a linguistics exam, the information in each of which could be filtered in two different ways, resulting in a story that—according to our judgment—should make a conditional pertaining to the story appear uncertain, and a story that—also according to

¹⁵ Actually, there is a difference in the tense of the conditionals used in the two stories. As far as we can see, this difference is immaterial. But of course this claim might be refuted by future empirical research.

our judgment—should make the same conditional appear certain. (See again the appendix for a description of these stories.) For the certain version of the Mr. Smith story, $N = 86$, for the uncertain version, $N = 97$; for the certain version of the story about the apple, $N = 87$, for the uncertain version, $N = 96$; and for the certain version of the linguistics exam story, $N = 95$, for the uncertain version, $N = 88$.

3.2. Results

A chi-square test was performed to test the hypothesis that one of the conditionals in the Gibbard story is a certain inferential conditional and the other an uncertain inferential conditional, given that “probably” is a good linguistic marker for distinguishing between the two types of conditionals. A moderately strong association was found between type of context (certain / uncertain) and assessment of effect of insertion of “probably,” $\chi^2(3, N = 183) = 24.9, p < .0001$; Cramér’s $V = .369, p < .0001$. Inspection of the cell frequencies showed that 55 (66 %) out of the 83 participants who had received the certain version of the story judged the corresponding conditional without “probably” to fit best into the context, whereas only 22 (27 %) of the same 83 participants judged the conditional with “probably” to fit better; 5 of the remaining 6 participants thought both conditionals fitted equally well; 1 participant thought neither fitted well. By contrast, 51 (51 %) out of 100 participants who received the uncertain version of the story judged the conditional with “probably” to fit best into the context, and only 30 (30 %) of the same group of participants judged the conditional without “probably” to fit best into the context; 12 participants thought that both conditionals fitted well; 7 thought neither fitted well.

We asked participants about their familiarity with the game of poker, and an analysis of the data of participants who indicated that they knew the game of poker well—38 of the 183 participants—revealed a stronger association than the one we found for the group as a whole. In particular, Cramér’s $V = .449$ for this group. However, the effect size is only marginally significant ($p = .053$).

We obtained similar results for DeRose’s story, although here a chi-square test revealed a *strong* association between type of context (certain / uncertain) and effect of “probably,” $\chi^2(3, N = 183) = 64.4, p < .0001$; Cramér’s $V = .59, p < .0001$. More specifically, in the certain context, 84 (88 %) out of 96 participants judged the conditional without “probably” to fit best; only 8 participants (8 %) judged the conditional with “probably” to fit best; 3 thought both fitted well; 1 thought neither fitted well. On the other hand, in the uncertain context, 26 (30 %) out of 87 participants judged the conditional without “probably” to fit best; 41 participants (47 %) judged the conditional with “probably” to fit best; 7 thought both fitted well; 13 thought neither fitted well.

As for the story about Mr. Smith, a chi-square test again found a strong association between type of context (certain / uncertain) and effect of “probably,” $\chi^2(3, N = 183) = 60.2, p < .0001$; Cramér’s $V = .57, p < .0001$. Cell frequencies for the certain version were: 64 (74 %) out of 86 participants thought that the conditional without marker

fitted best; 8 (9 %) thought the conditional with marker fitted best, 7 thought both fitted well; 7 thought neither fitted well. Cell frequencies for the uncertain version were: 20 (21 %) out of 97 thought the conditional without marker fitted best; 56 (58 %) thought the conditional with marker fitted best; 9 thought both fitted equally well; 12 thought neither fitted well.

As for the story about the apple, a chi-square test found a very strong association between type of context (certain / uncertain) and effect of “probably,” $\chi^2(3, N = 183) = 107.9, p < .0001$; Cramér’s $V = .77, p < .0001$. Cell frequencies for the certain version were: 65 (75 %) out of 87 participants thought that the conditional without marker fitted best; 11 (13 %) thought conditional with marker fitted best; 9 thought both fitted well; 2 thought neither fitted well. Cell frequencies for the uncertain version were: 5 (5 %) out of 96 participants thought the conditional without marker fitted best; 75 (78 %) thought the conditional with marker fitted best; 4 thought both fitted well; 12 thought neither fitted well.

As for the linguistics test story, a chi-square test found a strong association between type of context (certain / uncertain) and effect of “probably,” $\chi^2(3, N = 183) = 54, p < .0001$; Cramér’s $V = .543, p < .0001$. Cell frequencies for the certain version were: 50 (52 %) out of 95 participants thought the conditional without marker fitted best; 30 (32 %) thought conditional with marker fitted best; 9 thought both fitted equally well; and 6 thought neither fitted well. Cell frequencies for the uncertain version were: 6 (7 %) of 88 thought the conditional without marker fitted best; 54 (61 %) thought the conditional with marker fitted best; 4 thought both fitted equally well; and 24 thought neither fitted well.

3.3. Discussion

Except possibly for the apple story, the scenarios presented to the participants may have triggered some associations that we were not quite able to control. For instance, in the linguistics test story, 32 % of the participants decided that the conditional with “probably” fits best in the certain context while only 53 % of them thought the conditional without the marker fits best. Even though the effect of “probably” here is statistically significant, it is still curious that so many people chose the conditional with “probably.” The certainty of this scenario is based on rules that are of human origin. Participants may have their own experiences with universities, assignments and test, and some of them might have experienced that rules of this sort are sometimes violated.¹⁶ Hence, they may not be willing to conclude p on a basis of (roughly) “there is a rule that p ,” and they may still prefer “probably p ” in this kind of context.

The difference in the strength of association between Gibbard’s story and DeRose’s version of that story may be entirely due to the fact that the former involves the game of poker, the rules of which tend to strike outsiders as being relatively complex. Even

¹⁶It is well known that participants do not always fully believe what is given in vignettes and may, in view of their background beliefs, remain somewhat uncertain about what they are asked to suppose as being true. See Evans and Over [2004, Ch. 6] and Over, Douven, and Verbrugge [2013].

if no real knowledge of those rules is presupposed by the story, that itself is not immediately evident unless one is at least somewhat familiar with the game. By contrast, DeRose’s version poses no particular comprehension problems. For this reason, we take the responses to DeRose’s version to be more relevant to our hypothesis, and those responses revealed a strong association between type of context and effect of inserting “probably” in the relevant conditional. Specifically, inserting “probably” had a tendency to make the conditional appear better fitting in the uncertain context, but less fitting in the certain context. In view of the results of Krzyżanowska, Wenmackers, and Douven [2013], this finding is exactly what one would expect if the uncertain context invites an interpretation of the conditional as being uncertain and the certain context invites an interpretation of the conditional as being certain.

That, guided by our view on Gibbard’s story, we were able to design a number of scenarios that were meant to parallel Gibbard’s and DeRose’s and that turned out to elicit the same patterns in the responses, is further evidence for the correctness of our diagnosis.

4. Conclusion. Gibbard’s argument challenges truth-conditional accounts of conditionals. On the one hand, CNC appears an utterly plausible principle. On the other hand, the riverboat story features two conditionals, one of the form “If p , q ,” the other of the form “If p , not q ,” which in the context of the story strike one as being both true. According to Gibbard, it follows from his story that conditionals do not express propositions, and thus are not true or false; at least this follows, Gibbard thinks, unless one is willing to buy into a semantics that makes the communicative role of conditionals hard to understand. We have argued for a different conclusion by stating a semantics that relativizes the interpretation of conditionals to speakers’ background knowledge but not in a way that must leave us clueless about the meaning of a conditional whenever we are ignorant of the speaker’s background knowledge. As a consequence of our proposal, conditionals cannot be judged to be inconsistent with each other purely on the basis of their form. But this admission requires no more than a minor amendment in the formulation of CNC. Finally, it was seen that, in light of recent experimental results on linguistic markers, our analysis of Gibbard’s story has clear empirical content, part of which was investigated and borne out in the experiment that we conducted.¹⁷

Appendix

In addition to the two versions of Gibbard’s story presented earlier, we used the following materials in the experiment:

DeRose’s story, certain version: Jim and Steve are in the final round of a card game. They each draw one card from a deck of one hundred cards, numbered 1 through 100.

¹⁷We are grateful to Janneke Huitink, two anonymous referees, and the editors of the special issue for valuable comments. A version of this paper was presented at the 8th Barcelona Workshop on Conditionals. We thank the audience on that occasion for helpful questions and discussion.

Each of the players can look at his card, but not at the opponent's card. In this final round, it is up to the leader—in this case Jim—to decide whether to quit or play. If he decides to quit, he can keep the \$10,000 he won before this round. If he decides to play, both players have to show their cards. If Jim's card is higher than Steve's, his winnings will be doubled. But if he decides to play, and his card is the lower one, he will lose everything. In this final round, a friend of theirs who is present sees that Jim has drawn a card with the number 12 on it whereas Steve has drawn a card with the number 78 on it.

Conditionals: (1) "If Jim plays he will lose"; (2) "If Jim plays he will probably lose."

DeRose's story, uncertain version: Jim and Steve are in the final round of a card game. They each draw one card from a deck of one hundred cards, numbered 1 through 100. Each of the players can look at his card, but not at the opponent's card. In this final round, it is up to the leader—in this case Jim—to decide whether to quit or play. If he decides to quit, he can keep the \$10,000 he won before this round. If he decides to play, both players have to show their cards. If Jim's card is higher than Steve's, his winnings will be doubled. But if he decides to play, and his card is the lower one, he will lose everything. Before Jim made his final decision, his friend managed to see the number on Steve's card and made special hand signs to Jim to communicate the number to him.

Conditionals: (1) "If Jim plays he will win"; (2) "If Jim plays he will probably win."

Mr. Smith, certain version: Mr. Smith, a wealthy man, has two sons, Hank and Paul. Last week, he ordered his notary to draft a will. Mr. Smith is very fond of Hank but disappointed in Paul. The notary made the will, which mentions by mistake Paul instead of Hank as the main beneficiary of the Smith fortune. At this moment, Mr. Smith has an appointment at the notary's office to sign his will.

Conditionals: (1) "If Mr. Smith is not making any changes to his will, he is making Paul the main beneficiary of the Smith fortune"; (2) "If Mr. Smith is not making any changes to his will, he is probably making Paul the main beneficiary of the Smith fortune."

Mr. Smith, uncertain version: Mr. Smith, a wealthy man, has two sons, Hank and Paul. He is very fond of Hank but disappointed in Paul. In fact, on several occasions he has expressed his intention to make Hank the main beneficiary of his fortune. Last week, he ordered his notary to draft a will. At this moment, Mr. Smith has an appointment at the notary's office to sign his will.

Conditionals: (1) "If Mr. Smith is not making any changes to his will, he is making Hank the main beneficiary of the Smith fortune"; (2) "If Mr. Smith is not making any changes to his will, he is probably making Hank the main beneficiary of the Smith fortune."

Apple, certain version: The evil Mr. Saxon has poisoned an apple that he put on the table in the kitchen that you share with your house-mate Jack. The poison used by

Mr. Saxon is so effective that everyone who would eat the apple would get immediately sick. The apple is no longer on the table.

Conditionals: (1) “If Jack ate the apple, he is sick”; (2) “If Jack ate the apple, he is probably sick.”

Apple, uncertain version: Your house-mate Jack refused to join you for a pizza, saying that he feels sick and cannot eat anything at all. Half an hour later you notice that the apple that was lying on the table in your shared kitchen has disappeared.

Conditionals: (1) “If Jack ate the apple, he is well”; (2) “If Jack ate the apple, he is probably well.”

Linguistics exam, certain version: Students can get bonus points if they hand in the weekly exercises. In this case, if they get an A for the exam, their overall result is an A+, if they get an A- for the exam, they get an A for the course, and so on. If they don't hand in the exercises, they don't get any bonus points. Jim is a very talented student who has received an A+ for most of the courses he has taken. Yesterday Jim had a linguistics exam. The teacher has a look at Jim's test paper and he marks it with an A.

Conditionals: (1) “If Jim didn't hand in any exercises for the course, he won't get an A+ for it”; (2) “If Jim didn't hand in any exercises for the course, he probably won't get an A+ for it.”

Linguistics exam, uncertain version: Students can get bonus points if they hand in the weekly exercises. In this case, if they get an A for the exam, their overall result is an A+, if they get an A- for the exam, they get an A for the course, and so on. If they don't hand in the exercises, they don't get any bonus points. Jim is a very talented student who has received an A+ for most of the courses he has taken. It was often the case that if Jim thought he would be able to score an A+ even without handing in any exercises, he did not bother to do the exercises. Yesterday, Jim had a linguistics exam.

Conditionals: (1) “If Jim didn't hand in any exercises for the course, he will get an A+ for it”; (2) “If Jim didn't hand in any exercises for the course, he will probably get an A+ for it.”

References

- Adams, E. W. [1975] *The Logic of Conditionals*, Dordrecht: Reidel.
- Appiah, A. [1985] *Assertion and Conditionals*, Cambridge: Cambridge University Press.
- Bennett, J. [2003] *A Philosophical Guide to Conditionals*, Oxford: Clarendon Press.
- DeRose, K. [2010] “The Conditionals of Deliberation,” *Mind* 119:1–42.
- Douven, I. [2008] “The Evidential Support Theory of Conditionals,” *Synthese* 164:19–44.
- Douven, I. [2012] “The Lottery Paradox and the Pragmatics of Belief,” *Dialectica* 66:351–373.

- Douven, I. and Verbrugge, S. [2010] “The Adams Family,” *Cognition* 117:302–318.
- Edgington, D. [1995] “On Conditionals,” *Mind* 104:235–329.
- Evans, J. St. B. T. and Over, D. E. [2004] *If*, Oxford: Oxford University Press.
- Gabbay, D. and Woods, J. [2005] *The Reach of Abduction*, Amsterdam: Elsevier.
- Gibbard, A. [1981] “Two Recent Theories of Conditionals,” in W. L. Harper, R. Stalnaker, and G. Pearce (eds.) *Ifs*, Dordrecht: Reidel, pp. 211–247.
- Johnson-Laird, P. N. and Byrne, R. M. J. [2002] “Conditionals: A Theory of Meaning, Pragmatics, and Inference,” *Psychological Review* 109:646–678.
- Krzyżanowska, K. [2012] “Ambiguous Conditionals,” in P. Stalmaszczyk (ed.) *Philosophical and Formal Approaches to Linguistic Analysis*, Frankfurt: Ontos Verlag, pp. 315–332.
- Krzyżanowska, K., Wenmackers, S., and Douven, I. [2013] “Inferential Conditionals and Evidentiality,” manuscript.
- Kyburg, H. and Teng, C. [2001] *Uncertain Inference*, Cambridge: Cambridge University Press.
- Lewis, D. K. [1976] “Probabilities of Conditionals and Conditional Probabilities,” *Philosophical Review* 85:297–315.
- Lowe, E. J. [1996] “Conditional Probability and Conditional Beliefs,” *Mind* 105:603–615.
- Lycan, W. G. [2001] *Real Conditionals*, Oxford: Oxford University Press.
- McGee, V. [1985] “A Counterexample to Modus Ponens,” *Journal of Philosophy* 82:462–471.
- Over, D. E., Douven, I., and Verbrugge, S. [2013] “Scope Ambiguities and Conditionals,” *Thinking and Reasoning*, in press.
- Sanford, D. H. [1989] *If P, then Q*, London: Routledge.
- Stalnaker, R. [1975] “Indicative Conditionals,” *Philosophia* 5:269–286.
- Unterhuber, M. [2013] *Possible Worlds Semantics for Indicative and Counterfactual Conditionals?*, Frankfurt: Ontos Verlag.