

Indicatives at Stake*

Javier González de Prado Salas

jgonzlaezdeprado@gmail.com

IFILNOVA, Universidade Nova de Lisboa, Portugal

**This is a penultimate draft of an article to appear in Inquiry (DOI: 10.1080/0020174X.2018.1477617). Please refer to the published version.*

Abstract: Several authors have claimed that indicative conditionals are sensitive to the epistemic perspective of agents. According to this sort of view, the truth of an indicative conditional depends on the background evidence of some relevant agent or group of agents. In this paper I argue that the context-dependence of indicative conditionals goes beyond this. Indicative conditionals are not only sensitive to the evidence of agents, but also to contextual factors that determine what is inferable from such background evidence (plus the antecedent of the conditional). More specifically, my proposal is that when the inference associated with a conditional is defeasible, the truth of that conditional is sensitive to practical stakes (in a similar way that knowledge claims are sensitive to practical stakes).

Keywords: indicative conditionals; epistemic conditionals; defeasible reasoning; contextualism; stakes.

*For their comments and feedback, I thank Luke Burke, Lee Walters, Daniel Whiting, several anonymous referees and the audience at the 2016 meeting of the Spanish Society for Analytic Philosophy in Oviedo. This work has been supported by the Spanish Government research project FFI2014-57258-P ("Normative inferences and interferences in scientific research") and the Portuguese Foundation for Science and Technology project UID/FIL/00183/2013 ("Values in Argumentative Discourse").

1. Introduction

The topic of this paper is open indicative conditionals, that is conditionals in which the antecedent is treated as an open possibility. There is a close connection between such conditionals and reasoning. We typically use indicative conditionals to talk about what follows from the truth of some possibility – to show our endorsement of some inferential link. Think, for instance, of (1):

(1) If the water in the kettle is boiling, then it is hot

In general, those who accept (1) will be disposed to infer that the water in the kettle is hot from the fact that it is boiling.

In this paper I will discuss indicative conditionals associated with non-monotonic, defeasible inferences, i.e. inferences that may be made bad by the introduction of new, consistent premises. Consider, for instance, the inference represented by (2):

(2) If you press the brake pedal, the car will slow down

This inference is typically good, but it would become bad with the addition of the further information that the pedal is disconnected from the car's brakes. Defeasible reasoning is reasoning with incomplete information: a defeasible inference may be good (and non-enthymematic) even if there exist defeating possibilities compatible with the premises that would make the conclusion come out false (e.g. the possibility that the brake pedal is disconnected). Arguably, it is a context-sensitive matter whether a given defeating possibility may remain ignored or rather has to be explicitly ruled out by the premises of the inference (in order for it not to be enthymematic). My aim is to show that this context-sensitivity is reflected in the behavior of the conditionals associated

with defeasible inferences. In particular, it is plausible to think that practical factors (such as how high stakes are) contribute to setting the threshold of relevance under which defeating possibilities may be ignored; in other words, practical factors play a role in determining how much support is required for a given conclusion to be properly inferable from some premises. The thesis that I will explore here is that, accordingly, the truth value of conditionals linked with defeasible inferences also depends on such practical factors. Given that defeasible inferences are widespread in human reasoning and argumentative practices, this result illuminates an important dimension of our use of indicative conditionals.

This is the structure of the paper. In section 2, following Gillies (2004, 2009) and Weatherson (2009), I introduce an epistemic theory of indicative conditionals based on their connection with reasoning. Next, in section 3, I examine defeasible inferences and the conditionals associated with them. In section 4 I discuss what makes it the case that a defeating possibility is relevant enough and cannot be properly ignored. After that, in section 5, I argue that practical factors have a say in fixing the threshold of relevance for defeating possibilities. Before concluding, in section 6 I briefly consider the issue of which practical perspective is the one from which the threshold of relevance for a given conditional should be fixed.

2. Epistemic views of indicative conditionals

Ramsey (1929: 247) highlighted the connection between indicative conditionals and reasoning in the following famous remark:

‘If two people are arguing “If p , will q ?” and are both in doubt as to p , they are adding p hypothetically to their stock of knowledge, and arguing on that basis about q ’

One possible way of interpreting Ramsey’s remark is as pointing towards an epistemic theory of indicative conditionals, according to which the truth of the conditional ‘If p , then q ’ is related to the goodness of the piece of reasoning from p (and collateral premises belonging to the relevant evidential background) to q . So, ‘If p , then q ’ would be true iff q follows from p plus the background evidence available to the relevant agents. In this way, indicative conditionals can be said to represent the epistemic impact of adding the antecedent to some relevant epistemic perspective. As Gillies’ (2009: 601) claims, indicative conditionals are ‘global tests on what must epistemically be the case, given the antecedent’: when we assess an indicative conditional, we consider how our epistemic commitments would change if we update our informational state with the antecedent. Theories of conditionals with this epistemic flavor have been developed, among others, by Gillies (2004, 2009) and Weatherson (2009). Moreover, Gillies (2009) has shown how this sort of epistemic theory may be formulated within the framework of dynamic semantics (see also Veltman 1996; Starr 2014; Willer 2013). Even if the conclusions I will reach in this paper can be adapted to a dynamic approach, I will stick to a more standard truth-conditional framework.

Before going forward, let me clarify what I mean by inferential goodness. As I will understand the notion, an inference is good if its conclusion follows from the premises, that is to say, if there is a suitable inferential link between the premises and the conclusion, so that being entitled to endorse the premises would allow one to gain entitlement (via competent reasoning) to endorse the conclusion. In this way, a good inference can have premises that are false, and an agent may assess a certain inference

as good even if she disbelieves its premises and accordingly is not disposed to endorse the conclusion (see McHugh and Way 2018: 155). On this understanding, what matters for the goodness of an inference is not the truth of the premises, but whether the premises and conclusion are properly connected (in other words, inferential goodness would correspond to validity in classical deductive logic, rather than to soundness).

1

In order to flesh out the epistemic approach to conditionals I have sketched above, it is useful to characterize inferential connections in terms of relations of inclusion among possible worlds. On a first approximation (which I will revisit below), inferential goodness may be modelled as strict implication: a good inference would be such that all the worlds compatible with the premises belong to the set of worlds compatible with the conclusion. If this is combined with the epistemic theory of conditionals outlined above, we get a (dynamic) strict conditional analysis of indicatives (see Gillies 2009). According to this view, the conditional ‘If p , then q ’ is true iff q is true in all the worlds compatible with the informational state resultant from adding p to the relevant background evidence – that is, iff the worlds in which the consequent is true contain all the worlds compatible with the expanded informational state resulting from combining the antecedent with the background evidence.² More precisely, as

¹ For discussion of the connection between inferential goodness and inferential validity, see Dogramaci 2017. Note that, since I will be dealing with defeasible inferences, I will not assume that good inferences are always deductively valid (as we will see below, good defeasible inferences may fail to be truth-preserving across all possible worlds).

² Formulated in this way, the view is similar to the closest-worlds epistemic accounts proposed by Nolan (2003) and Williams (2008). Nolan adds the further condition that, among the worlds compatible with the background knowledge and the antecedent, one has to focus on those that are maximally similar to the actual one. Below we will see

Gillies (2009) notes, what is required is that the consequent is true at these worlds and in relation to a body of evidence enriched by the antecedent (this qualification is important when the consequent includes evidence-sensitive expressions, such as epistemic modals).³

Clearly, on this epistemic approach to conditionals, the truth value of indicative conditionals is sensitive to the background body of evidence. Whether the consequent of a conditional follows from the antecedent plus a given body of evidence (which would serve as a source of collateral premises) depends on what information is included in that body of evidence; changes in such collateral information will induce changes in the epistemic commitments resulting from updating that information with the antecedent of the conditional. Consider the following conditional:

(3) If Tom took the algebra exam, he failed.

This conditional seems to be true relative to the collateral information that every student who took the algebra exam failed. However, it is arguably false in relation to an alternative body of background evidence containing the information that Tom has passed all the exams he has taken (and not including the information that every student who took the algebra exam failed).⁴

that, when dealing with defeasible reasoning, epistemic theories have to introduce an analogous restriction: the consequent needs to be true only in all *relevant* worlds compatible with the evidence and the antecedent.

³ The view must be refined in order to deal with antecedents including epistemic modals (see Yalcin 2007). Given my purposes here, I can gloss over these complications.

⁴ Gibbard (1981) has famously discussed this sort of case, where seemingly incompatible conditionals are both true – in the same world –, each in relation to different background information.

This sort of sensitivity to collateral evidence is found as well in restrictor views of indicative conditionals, which are dominant in linguistic semantics (see Kratzer 2012). According to restrictor views, ‘if’ is not a connective, but rather marks a restriction on the domain of an operator. In particular, in the sort of epistemic conditional under examination, the if-clause would restrict the domain of a (sometimes covert) epistemic modal in the consequent. Under the influence of the if-clause, the domain of the modal operator in question would be restricted to those worlds compatible with the antecedent. When the consequent does not include an overt modal, a tacit one is posited (in most cases, this tacit modal is taken to be the epistemic necessity operator ‘must’). In this way, the restrictor view can be said to capture the epistemic dimension of Ramsey’s original proposal. In accordance with Ramsey’s remark, the restrictor view maintains that assessing an (epistemic) indicative conditional amounts to evaluating the epistemic repercussions of enriching the relevant body of evidence with the antecedent.⁵ On this view, indicative conditionals inherit the sensitivity to collateral evidence distinctive of modal operators.

An interesting question is which body of evidence should be considered in order to evaluate indicative conditionals. Is it the evidence possessed by the speaker? Or rather the evidence available to the subject assessing the conditional? Or perhaps it is constituted by the conversational common ground? I may remain neutral about these questions, since nothing of what I will say hangs on a particular answer to them (for discussion, see Weatherson 2009).

⁵ The restrictor view can be seen as a way of providing compositional semantics for the sort of evidence-relative epistemic theory I have sketched (Rothschild forthcoming). The differences between the restrictor view and other epistemic theories in which ‘if’ is treated as a connective (e.g. Gillies 2004) concern the compositional structure attributed to conditional sentences; such differences would become manifest at the level of truth conditions (if they do at all) only in special cases, for instance in some interactions between conditionals and overt modals. These sorts of cases will not be particularly

Most of the points I want to make could be made within the framework of the restrictor theory, from the point of view of dynamic semantics or even from the perspective of closest-worlds theories such as Nolan's (2003) and Williams' (2008). However, for the sake of simplicity I will focus on the simple epistemic theory that I have introduced at the beginning of this section.⁶ It is not my purpose in this paper to offer arguments favoring such epistemic theories of conditionals over alternatives such as the material conditional account (for a defense of epistemic theories, see Gillies 2004, 2009; Weatherson 2009). Rather, I want to explore an interesting form of context-dependence arising in these theories. More specifically, I will investigate the context-sensitivity associated with indicative conditionals codifying defeasible, non-monotonic inferences.

So far, I have assumed that, in order for an inference to be good, the premises (perhaps including collateral premises provided by background information) must strictly implicate the conclusion – the worlds compatible with the premises must belong to the worlds compatible with the conclusion. This is the case with monotonic inferences, like the one captured by the conditional 'If France is bigger than Monaco, then Monaco is smaller than France' or the inference represented by 'If John took the algebra exam, he failed (given that every student who took the algebra exam failed).' However, this assumption has to be relaxed in order to make room for non-monotonic inferences, which may count as good even if there are some possibilities compatible with the premises but incompatible with the conclusion. The view that I explore in this paper is that the truth of conditional associated with defeasible inferences requires only

relevant for my purposes here (for comparison between the two approaches, see Gillies 2010; Khoon 2011).

⁶I will also assume that conditionals have truth values, although the main conclusions of the paper could also be put forward in terms of acceptability conditions.

that all *relevant enough* possibilities compatible with the premises (the antecedent plus the evidential background) are compatible with the conclusion: possibilities below a certain threshold of relevance may be properly ignored. This (defeasible) strict conditional analysis can be formulated as follows:

Defeasible Strict Conditional: ‘If p , q ’ is true (relative to a body of evidence Σ and a threshold of relevance τ) iff all p -worlds compatible with Σ and whose relevance is above τ are worlds in which ‘ q ’ is true (in relation to a body of evidence Σ' resulting from updating Σ with p).

On this view the truth of an indicative conditionals does not only depend on the relevant background evidence, but also on what follows from such evidence plus the antecedent – which in turn depends on what possibilities are relevant and cannot be ignored. My aim is to examine what makes a given possibility *relevant enough*. I will argue that this is an occasion-sensitive matter; more specifically, I will claim that it is plausible to think that practical factors such as the conversational stakes play a role in determining the threshold of relevance under which a given defeating possibility may be ignored. If this is so, the truth of indicative conditionals will be sensitive to such practical issues.

In order to develop this proposal, in the next section I examine defeasible inferences in more detail.

3. Defeasible conditionals

The inferences associated with indicative conditionals are often non-monotonic (i.e. defeasible): it may be that the inference becomes bad after the introduction of a new premise compatible with the original ones (and with the background evidence providing collateral premises). When an inference is defeasible in this way, there are possible

worlds in which the premises are all true but the conclusion is not. Some examples are found in the inferences associated with the following conditionals (plus the background evidence indicated):

(4) If Mary's pet lays eggs, then it is not a mammal.

(5) If the clothes have been hanging outside, they will be dry by now.

(Available background evidence: It is sunny)

If Mary's pet happens to be a platypus, then the premises of the inference represented by (4) may be true despite its conclusion being false. Likewise, if your neighbor mischievously soaked the hanging clothes with her hose, then the clothes will not be dry – although it is actually a sunny day.

Assume, as a stipulated feature of the examples, that the bodies of background evidence relevant for assessing (4) and (5) do not rule out these defeating possibilities. So, the relevant background evidence in relation to (5) (say, the evidence of the speaker) establishes that it is sunny, but is silent about whether the clothes have been soaked by a mischievous neighbor. Thus, in both (4) and (5), there would be possibilities compatible with the antecedent and the available collateral evidence that make the consequent false. When such defeating possibilities obtain, the relevant inferences become bad.

On the view I will adopt here, an inference is good only if the premises rationally support the conclusion, that is, if the premises offer good reasons to endorse the conclusion (in the sense that, if the premises were true, they would constitute good factual reasons to accept the conclusion). In this way, an inference will be defeasible when the reasons offered by the premises can be defeated by further considerations, either because these defeating considerations provide stronger reasons in favor of an

incompatible conclusion or because they disable the reason-giving power of the premises.

Thus, if a defeater obtains, the original premises stop providing good enough reasons in favor of the conclusion, so that the inferential link connecting those premises and the conclusion becomes bad.⁷ On the epistemic theories I am considering, the truth value of a conditional is determined by the goodness of the relevant associated inference; therefore, when an (undefeated) defeater for some inference obtains, the conditional related to that inference is not true.

However, this does not mean that (4) and (5) are false in ordinary cases where defeating possibilities do not obtain. I may say something true by uttering (4) *even if the possibility that Mary's pet is a platypus has never crossed my mind and is not discarded by my evidence*. In the absence of defeaters, a defeasible piece of reasoning can be good despite the fact that its premises (including the collateral premises provided by the background evidence) do not rule out all possible defeaters. Otherwise, engaging in good reasoning would be too demanding: we would need to add enough collateral premises to rule out all defeating possibilities (and, in many cases, it seems that the list of possible defeaters will be indefinitely long). It is doubtful that agents will typically have enough evidence to discard all these possible defeaters. The point of defeasible reasoning is, precisely, to provide good reasons for a conclusion in cases where complete information is unavailable. Therefore, I will assume that when defeaters are actually absent, conditionals like (4) and (5) may be true – even if the relevant collateral evidence does not discard all possible defeaters.

⁷ Unless the defeater is itself defeated by further considerations. For the sake of simplicity, I will not dwell on this possibility here.

Most epistemic theories of indicative conditionals focus on cases where the inference connecting the proposition in the antecedent (plus the relevant evidence) and the proposition in the consequent is indefeasible (for instance, Gillies 2004: 587; Weatherson 2009: 339). However, the conditionals that we use in our ordinary speech are very often associated with defeasible inferences – which is not surprising, since our reasoning about most mundane topics is in general defeasible. It is plausible to expect, therefore, that important properties of conditionals will be obscured by the simplifying assumption that the inferences associated with them are monotonic.

Actually, the failure of strengthening the antecedent, one of the trademark features of indicative conditionals, has its roots in the defeasibility of the inferences associated with the conditionals in question. Strengthening the antecedent of a conditional amounts to introducing a new premise in the inference associated with it. If this inference is non-monotonic, the addition of the new premise may make it become bad – so that the related conditional turns out to be false. Think of this strengthening of (4):

(6) If Mary's pet lays eggs and it is a platypus, then it is not a mammal.

(6) is clearly false, since the claim in the consequent does not follow from the antecedent (plus any background of evidence). By contrast, if the inference associated with a conditional is indefeasible, then strengthening of the antecedent will never fail. Consider (7):

(7) If n is an even number, then n is a multiple of 2.

No matter what further claim is added to the antecedent of (7) (as long as it remains consistent), the conditional will be true.

In order to account for the failure of strengthening the antecedent, Gillies (2004: 607-609) relaxes momentarily the assumption that the inferences associated with conditionals are indefeasible; instead, such inferences are treated as indefeasible only within a restricted domain constituted by those possibilities that are relevant in the context. Other, non-relevant possibilities may be properly ignored. Appealing to Lewis' (1996) rules for ignoring possibilities, Gillies argues that the claim introduced in the strengthened antecedent counts automatically as a relevant possibility. More specifically, Gillies thinks that this is so because of Lewis' Rule of Attention, according to which possibilities to which one pays attention cannot be properly ignored (Lewis 1996: 559). The idea is that, when some possibility is introduced in the expanded antecedent of the conditional, we are paying attention to it and therefore counts as relevant.

Although I think that Gillies' analysis of strengthening the antecedent is in the right direction, I do not think that appealing to the Rule of Attention is the best way to go here. I take it that it is controversial whether some possibility automatically becomes relevant just because we pay attention to it – or because it somehow becomes conversationally salient (for discussion, see McKenna 2014; Ichikawa 2011a; Williams 2001a; Blome-Tillmann 2009).⁸ One does not need to grant that this is so in order to deal satisfactorily with the failure of strengthening the antecedent. When some proposition p is introduced in the antecedent of an indicative conditional, the possibility that p is not merely made salient. As Ramsey claims, p is hypothetically added to our

⁸ Note, however, that Lewis is right that when we pay attention to some possibility we are not actually ignoring it, even if perhaps we could have properly ignored it (see Lewis 1996: 559). I am disputing the Attention Rule as a rule for determining the relevance of possibilities, not as a rule saying when something is *actually* (properly) ignored.

‘stock of knowledge’: p would become an assumption in our piece of reasoning – we examine what follows from some body of evidence after p is incorporated into it. Thus, within the context of such a piece of conditional reasoning, one treats p as being the case. Very plausibly, a possibility that is assumed to be actual should count, within the scope of such an assumption, as relevant – in line with the spirit of Lewis’ Rule of Actuality (1996: 554).

Strengthening the antecedent of a conditional, therefore, is an explicit mechanism for making relevant some possibility that perhaps was previously ignored. But, are there other ways in which possibilities may become relevant? What factors determine whether some possibility may remain ignored or rather needs to be considered when assessing a given conditional? This is the question that I explore in the rest of the paper. I will argue that relevance of possible defeaters is a contextual, occasion-sensitive issue and, more specifically, that practical matters determine how relevant some possibility has to be in order to count as relevant enough.

4. When is a defeating possibility relevant?

For most defeasible inferences, there are a large number of possible defeaters (often, indefinitely many). For instance, the inference represented by (5) will be defeated not only if your neighbor soaks the clothes with her hose, but also if some kids are playing with their water guns around the hanging clothes, if the garden sprinklers have activated by mistake and so on.

I will assume that, in general, agents engaging in defeasible reasoning are not in a position to rule out explicitly all possible defeaters, and that therefore there is a distinction between defeating possibilities that are discarded by the agent’s evidence

and defeating possibilities that are simply ignored. One could perhaps think that defeasible inferences are actually enthymemes, underlain by hidden premises excluding the presence of defeaters. This, however, is problematic. More often than not, our evidence does not rule out the presence of all defeating possibilities. Thus, if our reasoning were always supported by a (perhaps indefinitely large) number of hidden anti-defeater premises, our reasoning would only rarely give us reasons to endorse the conclusion – since, in general, we would not have reasons to accept many of the hidden anti-defeater premises (see Horty 2012: 85-90). I will assume as a given fact that we often succeed in acquiring reasons to endorse some conclusion by engaging in defeasible reasoning. In order for this sort of defeasible reasoning to leave the ground, it has to be possible for an inference to be good (rather than enthymematic) even if its premises do not rule out all possible defeaters for the goodness of such an inference. Accordingly, an indicative conditional may be true even if its antecedent plus the relevant background evidence do not exclude all possibilities that would defeat the associated inference.

Perhaps some may want to insist that we often have enough evidence to rule out most possible defeating possibilities, or even to endorse a material conditional connecting the premises and conclusion of the inference (for arguments in this direction, see Bird 2004). However, if this were so, it would be difficult to explain why indicatives in which the antecedent is strengthened with the introduction of a defeating condition are still read as (false) *indicative* conditionals, rather than as *counterfactuals* (see, for instance, (6) above). After all, indicatives are customarily taken to presuppose that the antecedent is compatible with the background evidence. So, if the defeating possibility added to the antecedent were ruled out by our evidence, the corresponding enriched conditional should be read as a counterfactual and not as an indicative. It

seems, therefore, that such defeating possibilities are disregarded or ignored, rather than ruled out by the agent's background evidence.

So, not all defeating possibilities have to be explicitly considered when engaging in defeasible reasoning; some such possibilities may be (properly) ignored. In this way, whether some conclusion can be inferred from a given premise depends on the collateral premises available *and also on which defeating possibilities may be properly ignored*. But, what determines whether a given possibility can remain unconsidered or rather has to be ruled out by the premises of the inference? I will say that a defeating possibility requires explicit consideration when it is *relevant enough*. Then, the question can be reformulated as the question about when, and by virtue of what, some defeating possibility counts as relevant enough.

As suggested in the discussion of strengthening the antecedent in the previous section, one could resort at this juncture to Lewis' rules for properly ignoring possibilities (Lewis 1996). Lewis's rules are concerned with knowledge ascriptions, but, arguably, analogous rules should be in play in relation to the assessment of the goodness of inferences. One of the typical ways of getting to know some proposition is to infer it correctly from known premises. Thus, it is to be expected that there will be a close connection between the possibilities that may be ignored when assessing whether one knows the conclusion of some inference and the possibilities that may be ignored when assessing the goodness of the inference that led to such purported knowledge.

I think that the rules described by Lewis certainly point in the right direction. However, I do not want to commit myself here to the details of Lewis' proposal (as I have already mentioned, some of his rules, in particular the Rule of Attention, are controversial). For my purposes here, I do not need to specify an exhaustive set of rules or criteria that fully determine when a possibility is relevant enough. It will be enough

to assume that not all possibilities are equally relevant, in relation to a given inference and a given context, and that in general some less relevant possibilities may remain properly unconsidered. So, instead of putting forward a specific, detailed account of what makes a possibility relevant, I will make some general remarks, in an exploratory spirit, with the aim of providing some illumination on the notion of relevance I have in mind (for further discussion, see González de Prado Salas 2017).

A first plausible idea is that defeating possibilities may become relevant because the participants of the conversation explicitly leave them open, expanding the common ground so as to include them (see Blome-Tilmann 2009); arguably, this is what happens when the assertion of a sentence involving the epistemic modal ‘might’ (for example ‘Your neighbor might have soaked the clothes with the hose’) is accepted and the common ground updated accordingly (Williams 2008).

Moreover, it seems that a defeating possibility will be relevant enough if there exist strong enough reasons to consider it. Now, whether there exist good reasons to consider the possible presence of a certain defeater depends, arguably, on features of the circumstances of inference. In particular, it is plausible to think that it will depend among other things on whether the defeating possibility is abnormal in the sort of situation targeted by the inference (see González de Prado Salas 2017). If the defeater is typical in the environment the inference is about, then it will be relevant and will have to be explicitly considered. By contrast, defeating possibilities that are abnormal or atypical in that environment may be ignored, insofar as they are unlikely to obtain and their presence would require special explanation (see Smith 2016).⁹

⁹Note that, when one’s evidence is purely statistical, explicitly assigning even a very low probability to a defeating possibility may be enough to make it relevant enough. Arguably, this is what happens in lottery examples.

For instance, (4) represents a good inference if Mary is in an ordinary environment, where it is unheard of that someone's pet is a platypus; in such ordinary scenarios, that possibility may be ignored – and therefore the conditional (4) will be true even if the collateral evidence says nothing about Mary's pet not being a platypus. However, if platypuses were common pets in Mary's surroundings, then such a possibility would become relevant and it would be inappropriate to ignore it: the inference associated with (4) would be enthymematic unless there were a further premise ruling out the possibility that Mary's pet is indeed a platypus. Thus, in this second scenario the conditional (4) would be false if the relevant body of evidence did not include the information that Mary's pet is not a platypus.

It may well be that the relevance of a defeating possibility does not depend only on what the typical environmental conditions are, but also on further contextual features such as the presuppositions of the conversation, the interests of agents and the sort of inquiry they are engaged in (see Williams 2001b). Be this as it may, it seems that occasion-sensitive factors such as environmental typicality do contribute to deciding whether a defeating possibility can remain ignored or has to be explicitly considered by the premises of the inference (on pain of being enthymematic). In this way, whether some conclusion is (defeasibly) inferable from some premises is an occasion-sensitive matter. As a result, it will also be an occasion-sensitive matter whether an indicative conditional associated with a defeasible inference is true – given that, on the epistemic view I am working with, the truth value of indicative conditionals is linked to the goodness of the inference associated with them. Thus, the truth value of indicative conditionals is not only relativized to a background body of evidence, but also to those occasion-sensitive factors that determine what defeating possibilities are relevant enough (e.g. the typical environmental conditions).

Jonathan Ichikawa (2011b) and Karen Lewis (2016) have both put forward contextualist accounts of counterfactuals according to which it is a context-sensitive question what possible worlds are relevant for assessing a given counterfactual. Lewis' (2016) main motivation for endorsing a contextualist approach is to avoid counterfactual skepticism – that is, to avoid the result that intuitively acceptable counterfactuals are actually false due to the existence of fancy possibilities compatible with the antecedent and which seem to figure among the closest possible worlds. Lewis' solution is to argue that such fancy possibilities are not relevant in ordinary contexts (she holds that relevance is mainly determined by the purposes of the conversation (Lewis 2016: 293)).

Ichikawa's proposal, in turn, is motivated by (reverse) Sobel sequences and their relations to analogous puzzles about knowledge ascriptions. In order to deal with such cases, Ichikawa proposes the following account of counterfactuals, mimicking Lewis (1996) account of knowledge:

If A were the case, C would be the case is true just in case all of the A possibilities are C possibilities (Psst!—except those possibilities we're properly ignoring). (Ichikawa 2011b: 296)

Although there are similarities between the spirit of Lewis' and Ichikawa's proposals and mine, the differences are also clear. First, I am focusing on indicative conditionals, whereas they are interested in counterfactuals. Moreover, my proposal is not specifically introduced as a solution to a specific puzzle or difficult case, but is rather supported by general observations about defeasible inferences and their connections (via Ramsey's remark) with indicative conditionals – although, of course, this does not mean that the proposal I am presenting cannot throw new light on several problematic cases, such as the failure of strengthening the antecedent, (reverse) Sobel sequences or

skeptical challenges to our inferential practices and our use of indicatives. At any rate, I take it that it is interesting to have a principled, unified picture of the relation between the defeasibility of inferential practices and the context-sensitive behavior of indicative conditionals.

Actually, the context-sensitive nature of defeasible reasoning that I have discussed may be of help to understand the context-sensitivity not only of conditionals, but also of knowledge ascriptions, given the close relation between reasoning and knowledge acquisition. In the next section, I explore a particularly interesting parallel between conditionals and knowledge ascriptions, which has not received too much attention so far. It is a popular idea in the recent literature that attributions of knowledge are sensitive to practical factors, because such factors set, on each occasion, the threshold of evidence required to count as knowing (see, among others Stanley 2005; Schroeder 2012). I will argue that something similar happens with conditionals. Practical factors have a say in determining the threshold of relevance for a defeating possibility to merit explicit consideration. In this way, practical factors can have an influence on the truth value of indicative conditionals.

5. Practical stakes and conditionals

I have argued that whether a defeating possibility needs to be considered when assessing the goodness of an inference depends on its relevance. But I have not addressed yet the question of how relevant is *relevant enough*. Which is the threshold of relevance under which a defeating possibility may be disregarded, without making the inference enthymematic and undermining the truth of the associated conditional? How is this threshold settled?

One option is to think that there is a fixed, invariant threshold that applies in all cases. I will explore an alternative possibility, according to which the threshold may change from context to context – in particular, it may change due to practical factors such as how high the stakes of the conversation are.¹⁰ The idea that there may be a variable relevance threshold becomes natural once one abandons the view that an inference is good only if all defeating possibilities are ruled out by its premises.

The analogous view about knowledge ascriptions (that is, the view that the correctness of knowledge ascriptions depends on practical stakes) has been widely discussed in recent years (see DeRose 1992; Cohen 1999; Stanley 2005; Fantl and McGrath 2009; Schroeder 2012).¹¹ On this view, when stakes are high and the costs of

¹⁰ Alternatively, one could think that what happens is not so much that practical factors shift the relevance threshold, but rather that such factors induce global changes in how relevant defeating possibilities are. For my purposes here, these two options can be taken to be equivalent. What matters is whether practical issues can alter the relative relevance of defeating possibilities with respect to the operative threshold, and not so much whether what changes is the threshold or the level of ‘absolute’ relevance of the possibilities. The important point is that, on both approaches, practical factors would have an impact on how many possibilities get to be above the relevance threshold. For ease of exposition, I will assume that it is the threshold what is lowered or raised.

¹¹ The thesis that the truth value of knowledge attributions is sensitive to practical stakes is popular, but not uncontroversial: for criticism see, among others, Buckwalter and Schaffer (2015) and Turri (2017). Those who are reluctant to accept the thesis in relation to knowledge attributions will probably also have reservations about the analogous position about indicative conditionals explored here – although it would be interesting to examine in detail how the considerations adduced against the thesis about knowledge ascriptions affect the corresponding thesis about conditionals. Due to space constraints, discussion of these critical views will have to wait for another occasion. Note, at any rate, that some of the alleged empirical evidence against the thesis of the pragmatic sensitivity of knowledge attributions would not affect relativist views,

error increase, it becomes harder to count as knowing – the threshold of evidential support required for being correctly attributed knowledge goes up. This phenomenon is often illustrated by means of the famous bank cases (DeRose 1992: 913; Stanley: 3-5):

Low stakes bank: Hannah and Sarah plan to go to the bank today, Friday, to deposit their paychecks. However, they see that the queue is very long. They wonder whether they may come back tomorrow Saturday, or whether it will be closed. Hannah has checked that the bank’s website says that it closes only on Sundays, so she claims to know that the bank will be open tomorrow Saturday. It is not very important that Hannah and Sarah deposit their paychecks by Saturday (they have no impending bills).

High stakes bank: Same as above, but in this case it is crucial that Sarah and Hannah deposit their paychecks by Saturday (their mortgage payment is due by then).

According to the orthodox interpretation (at least among philosophers) of these cases, in the first, low stakes situation it is correct for Hannah to claim to know that the bank opens on Saturday, relying on the evidence she got by checking the bank’s website. By contrast, in the second, high stakes situation (where the costs of failing to go the bank before Sunday are very high), Hannah’s evidential state would not be strong enough to

according to which the relevant stakes are those of the context of assessment (such as MacFarlane’s 2014: ch. 8). The reason for this is that the experiments considered typically involve changes in the stakes of the putative knower or of the attributor of knowledge, rather than the stakes of the assessor of the attribution (i.e. the experimental subject). In section 6 I identify the analogous relativist view for conditionals.

make her count as knowing that the bank will open on Saturday: she would need further evidence, in particular she would need to be able to rule out undermining possibilities such as the bank's website not being properly updated. Authors like Fantl and McGrath (2009) motivate this view by appeal to the link between knowledge and rational action. The idea is that it is always rational to act on a known proposition (for criticism, see Brown 2008). Arguably, in the high stakes scenario it is not rational for Hannah to act as she would act on the assumption that the bank will open on Saturday. Therefore, she cannot count as knowing that this is the case.

Consider now Hannah's reasoning. She concludes that the bank will be open on Saturday, since it says so in its website. Is Hannah's piece of reasoning good? Is she entitled to endorse such a conclusion, given her evidence, or does she need to rule out defeating possibilities such as the possibility that the bank's website is not updated? I want to suggest that it may be that such a defeating possibility can remain ignored in the low stakes case, while requiring explicit consideration in more demanding scenarios. Note that what changes is not Hannah's evidential state – in the sense that she does not acquire additional information concerning the likelihood of the bank's website not being accurate.¹² Rather, what happens is that when the costs of error raise, Hannah may need to considered explicitly low likelihood possibilities in order to be in a position to infer correctly that the bank will be open on Saturday. Given the costs she faces, she needs to be extra careful when drawing conclusions from her premises. Thus, I am putting forward the thesis that inferential goodness is sensitive to practical factors:

¹² Arguably, Hannah's evidential state could be said to change in the sense that in the high stakes situation her evidence allows her to infer fewer conclusions than in the low stakes scenario.

Inferential Pragmatic Sensitivity: practical stakes contribute to determining how much inferential support is required for an agent to become rationally permitted to endorse a certain conclusion.

It should be stressed that the view I am presenting finds support in the same sort of considerations that motivate the thesis that knowledge ascriptions are sensitive to practical factors. In particular, there is a clear link between the rationality of one's actions and the goodness of the inferences that led to such actions (plausibly, good reasoning preserves rational permissibility). Moreover, both views are underpinned by similar anti-skeptical intuitions: if it is granted that one can have knowledge that is not supported deductively by one's evidence, it seems that it should be granted more generally that there can be good, non-enthymematic defeasible inferences. If by relying on good defeasible reasoning from known premises one may acquire non-deductive knowledge, and ascriptions of non-deductive knowledge are sensitive to practical factors, then it is natural to think that the goodness of pieces of defeasible reasoning is also sensitive to practical facts.

A natural way of modelling this type of influence of practical factors on inferential goodness is by claiming that increments in the costs of inferential errors lower the threshold of relevance below which a defeating possibility may be ignored. Accordingly, in demanding situations, fewer things can be properly ignored. Thus, whether some conclusion can be said to follow from certain premises would depend on practical factors, because such factors determine how relevant a defeating possibility has to be in order to require explicit consideration by the premises of the inference. On a view of conditionals inspired by Ramsey's remark, this means that the correctness of an indicative conditional may be affected by practical factors. The same conditional, in

relation to the same background evidence, may be true in a low stakes situation and false in a high stakes case. In a high stakes situation, the support received by the conclusion of the inference has to be stronger in order for the associated conditional to be acceptable. Under the modelling assumption that what is shifted when practical stakes change is the operative relevance threshold, we reach the following result:

Pragmatic Sensitivity of Conditionals: changes in practical stakes may shift the relevance threshold τ determining the set of worlds that can be properly ignored when evaluating the truth value of a conditional ‘If p, q’.

Take, for instance, (8) below:

- (8) If the bank’s website says that the bank closes only on Sundays, then it will open tomorrow Saturday

The view I am putting forward is that (8), relative to a body of background evidence not ruling out that the bank’s website is inaccurate, may be true in a relaxed situation – where unlikely possibilities may be disregarded – and false in a high stakes context – where one is required to check further defeating possibilities before being entitled to jump to conclusions.

Ichikawa (2011b) considers briefly, in a footnote, whether counterfactuals are sensitive to practical stakes in an analogous way to knowledge attributions. He seems inclined to think that this sensitivity to stakes is weaker in the case of counterfactuals, and one of the considerations he adduces in order to explain this is that, on his view, ‘truths about counterfactual situations are much less likely to be of great practical

importance in deliberation' (2011b: 300 fn. 13).¹³ I do not think that the same can be said about indicative conditionals. Such conditionals have an important function in our deliberations about what to believe and what to do. One of the main uses of indicative conditionals is to examine and discuss what follows from our evidence, what view of things we should adopt given what we know. This use of indicative conditionals has obvious consequences for our decision-making. In particular, conditionals play a crucial role in paradigmatically high stakes scenarios, such as medical practice or legal contexts. In this way, whether a given conditional is accepted may determine the verdict of a jury or the treatment recommended by a doctor.¹⁴ In this sort of case, the costs of endorsing false conditionals may be high.

As an example, imagine that Tom's office had the lights on and a silhouette could be seen through its window yesterday at 9 p.m. Consider first a relaxed context, in which Tom's colleague Alice wonders whether Tom finally got to leave the office before 9 p.m. (say, she wants to know whether Tom managed to go to a concert that started at 9 p.m.). Given that Tom does not share his office with anyone, it seems that, in this low stakes context, Alice may infer that Tom was still in his office last evening at 9 – even if she has no evidence ruling out the unlikely possibility that someone broke in Tom's office and switched the lights on. Thus, in this context (9) will be taken to be true:

(9) If yesterday at 9 p.m. Tom's office had the lights on and a silhouette could be seen through the window, then Tom was in his office at that time.

¹³ Ichikawa (2011b: 300 fn. 13) suggests that high stakes do not shift standards for knowledge, but rather make agents pay attention to more possibilities, which thereby are not ignored anymore.

¹⁴ For what is worth, I also think that counterfactuals often have practical importance – for instance for attributions of moral and legal responsibility.

Consider now a high stakes scenario, say a police investigation in which it is crucial to determine who was in the company's building yesterday at 9 p.m. (a murder was committed in the building at that time). In this demanding context, the possibility that someone broke into Tom's office may need to be considered explicitly, so that the conditional will not be accepted unless such a possibility is ruled out by the evidential background (or by an expanded antecedent).

It seems plausible, thus, to think that in demanding contexts, where the costs of inferential errors are high, stronger premises are required in order to be entitled to reach inferentially some conclusion. As a consequence, changes in conversational stakes may bring changes in the truth value of indicative conditionals (at least, those conditionals associated with defeasible inferences).

Note that a similar claim can be advanced within the framework of non-truth-conditional accounts of indicatives (Adams 1975; Edgington 1995; Bennett 2003; Gibbard 1981). Even if according to these accounts indicative conditionals do not have truth values, they will arguably still have acceptability conditions: the assertion of an indicative conditional will only be correct in certain circumstances. Assume, for instance, that a conditional is acceptable only when the conditional probability of the consequent given the antecedent is sufficiently high. How high does this conditional probability need to be? On the view I am exploring, the answer would depend on practical factors: in demanding contexts, a higher probability of the consequent given the antecedent would be required in order for the associated indicative to be acceptable.

6. Which context?

I have suggested that the relevance threshold under which a defeating possibility may be ignored is sensitive to practical matters such as how high stakes are. However, I have not discussed yet whose practical perspective is the one that sets this threshold. Should the threshold be fixed from the perspective of the context of utterance, from the context of assessment or from some other practical point of view?

Arguably, these questions are related to analogous ones about the context-sensitivity of the threshold for knowledge. As I have explained in the previous section, it is a popular view that the correctness of knowledge attributions is sensitive to practical factors. It is a matter of debate, however, whether the practical facts relevant for this sort of attribution are facts pertaining to the subject's situation, or rather to the context of utterance or assessment of the sentence attributing knowledge (for a summary of this debate, see MacFarlane 2014: ch. 8). It is to be expected that one's stance on this debate will inform one's views about the context-sensitivity of the threshold of relevance for indicative conditionals – given the way in which our judgments about whether a subject has knowledge are connected with our judgments about whether the subject may correctly infer a given conclusion from some body of evidence.

Interestingly, however, one of the possible positions in the debate about knowledge attributions, subject-sensitive invariantism, does not seem to be available in relation to conditionals. According to subject-sensitive invariantism, the relevant practical facts are those related to the subject to whom knowledge is attributed, who may be different from the speaker making the attribution (see Hawthorne 2004; Stanley 2005). Yet when a speaker utters an indicative conditional, she is not assessing the epistemic state of some third party, but rather the epistemic position of those sharing her evidential background – more specifically, she is considering what is inferable (adding

the antecedent as a premise) for those in her conversational context.¹⁵ In this way, the relevant subject in relation to conditionals is the speaker herself (and, arguably, her conversational partners). Thus, those who favor subject-sensitive invariantism about knowledge attributions should take it that the relevant practical factors for assessing conditionals are those of the context of the speaker, thereby embracing a contextualist position about indicative conditionals.

On the other hand, relativist views about knowledge attributions (MacFarlane 2014: 187-190) would be mapped onto corresponding relativist positions about conditionals, according to which the practical facts that fix the relevance threshold are those of the context from which the conditional is assessed. I remain neutral here about whether a contextualist or a relativist account of the relevance threshold for indicative conditionals is to be preferred. I do not adjudicate either between *indexicalist* positions, according to which the proposition expressed by ‘If p, then q’ changes with changes of the relevance threshold, and *non-indexicalist* views, according to which the proposition remains the same and what varies are the circumstances with respect to which it is to be evaluated (I take this terminology from Weatherson 2009; see as well MacFarlane 2014: 78-92). I also leave for another occasion the discussion of how the context-sensitivity stemming from the determination of the relevance threshold interacts with the context-sensitivity associated with the choice of background evidence (this last form of context-sensitivity is examined by Weatherson 2009).

¹⁵ Arguably, the perspective-dependence of indicative conditionals is closer to that of the epistemic modal ‘must’ (which is not associated with a specific agent) than to that of ‘know’. It would be interesting to study in detail the sensitivity of ‘must’ to practical factors. I leave this issue for another occasion.

7. Conclusions

Indicative conditionals associated with non-monotonic inferences inherit the context-dependence characteristic of defeasible reasoning. The study of the properties of defeasible reasoning reveals the general, widespread way in which indicative conditionals are sensitive to the dynamics of discourse and to contextual factors – a sensitivity that explains (but goes beyond) distinctive features of the behavior of conditionals such as the failure of strengthening the antecedent. The truth value of indicative conditionals does not only depend on a background body of evidence but also on what defeating possibilities may be ignored on the occasion of use – since this determines what can be inferred from the background evidence plus the antecedent. In particular, I have argued that, plausibly, practical facts influence the truth value of indicative conditionals, by fixing the relevance threshold under which defeating possibilities may remain ignored without making the inference associated with the conditional enthymematic.

References

- Adams, E. (1975). *The Logic of Conditionals*. Dordrecht: D. Reidel.
- Bennett, J. (2003). *A philosophical guide to conditionals*. Oxford: Oxford University Press.
- Bird, A. (2005). Abductive knowledge and Holmesian inference. *Oxford studies in epistemology 1*, 1-31.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118(470), 241-294.
- Brown, J. (2008). Knowledge and practical reason. *Philosophy Compass*, 3(6), 1135-1152.

- Buckwalter, W. and Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs*, 49(2), 201-234.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Noûs*, 33(s13), 57-89.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and phenomenological research*, 52(4), 913-929.
- Dogramaci, S. (2017). Why Is a Valid Inference a Good Inference?. *Philosophy and Phenomenological Research*, 94(1), 61-96.
- Edgington, D. (1995). On conditionals. *Mind*, 104(414), 235-329.
- Fantl, J. & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Gibbard, A. (1981). Two recent theories of conditionals. In W.L. Harper, R. Stalnaker & G. Pearce (eds.), *Ifs*. Dordrecht: D. Reidel, 211–248.
- Gillies, A.S. (2004). Epistemic conditionals and conditional epistemics. *Noûs*, 38(4), 585-616.
- Gillies, A.S. (2009). On truth-conditions for if (but not quite only if). *Philosophical Review*, 118(3), 325-349.
- Gillies, A.S. (2010). Iffiness. *Semantics and Pragmatics*, 3, 4-1.
- González de Prado Salas, J. (2017). Defeasibility and Inferential Particularism. *International Journal of Philosophical Studies*, 25(1), 80–98.
- Horty, J.F. (2012). *Reasons as defaults*. Oxford: Oxford University Press.
- Ichikawa, J. (2011a). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155(3), 383-398.
- Ichikawa, J. (2011b). Quantifiers, knowledge, and counterfactuals. *Philosophy and Phenomenological Research*, 82(2), 287-313.

- Khoo, J. (2011). Operators or restrictors? A reply to Gillies. *Semantics and Pragmatics*, 4, 4-1.
- Kratzer, A. (2012). *Modals and conditionals*. Oxford: Oxford University Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian journal of Philosophy*, 74(4), 549-567.
- Lewis, K.S. (2016). Elusive Counterfactuals. *Noûs*, 50, 286-313.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- McHugh, C., & Way, J. (2018). What is Good Reasoning?. *Philosophy and Phenomenological Research*, 99, 153-174 .
- McKenna, R. (2014). Normative scorekeeping. *Synthese*, 191(3), 607-625.
- Nolan, D. (2003). Defending a Possible-Worlds Account of Indicative Conditionals. *Philosophical Studies*, 116, 215-269.
- Ramsey, F.P. 1929. General propositions and causality. In D. H. Mellor (Ed.) *Philosophical Papers*. Cambridge: Cambridge University Press, 1990, 145– 163.
- Rothschild, D. (forthcoming). A Note on Conditionals and Restrictors. In J. Hawthorne and L. Walters (eds.), *Conditionals, Probability and Paradox: Themes from the Philosophy of Dorothy Edgington*. Oxford: Oxford University Press.
- Schroeder, M. (2012). Stakes, withholding, and pragmatic encroachment on knowledge. *Philosophical Studies*, 160(2), 265-285.
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford: Oxford University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.

- Starr, W.B. (2014). A uniform theory of conditionals. *Journal of Philosophical Logic*, 43(6), 1019-1064.
- Turri, J. (2017). Epistemic Contextualism: An Idle Hypothesis. *Australasian Journal of Philosophy*, 95(1), 141–156.
- Veltman, F. (1996). Defaults in update semantics. *Journal of philosophical logic*, 25(3), 221-261.
- Weatherson, B. (2009). Conditionals and indexical relativism. *Synthese*, 166(2), 333-357.
- Williams, J.R.G. (2008). Conversation and conditionals. *Philosophical Studies*, 138(2), 211-223.
- Williams, M. (2001a). Contextualism, externalism and epistemic standards. *Philosophical Studies*, 103(1), 1-23.
- Williams, M. (2001b). *Problems of knowledge: A critical introduction to epistemology*. Oxford: Oxford University Press.
- Willer, M. (2013). Indicative scorekeeping. *Proceedings of the 19th Amsterdam Colloquium*, 249-256.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464), 983-1026.