

# Defeasibility and Inferential Particularism

Javier González de Prado Salas

Universidad Autónoma de Madrid, Spain

jgonzalezdeprado@gmail.com

*\*This is a penultimate draft of an article to appear in the International Journal of Philosophical Studies. Please refer to the published version.*

**Abstract:** In this paper I argue that defeasible (nonmonotonic) inferences are occasion-sensitive: the inferential connections of a given claim (plus collateral premises) depend on features of the circumstances surrounding the occasion of inference (including typical environmental conditions, and also pragmatic and contextual factors, such as the information available to agents and how high stakes are). More specifically, it is an occasion-sensitive matter which possible defeaters have to be considered explicitly by the premises of an inference and which possible defeaters may remain unconsidered, without making the inference enthymematic. As a result, a largely unexplored form of occasion-sensitivity arises in inferentialist theories of content that appeal to defeasible inferences.

## 1. Introduction

Defeasible (nonmonotonic) inferences are inferences whose goodness can be defeated by the addition of new (consistent) premises – i.e. defeaters. Such inferences are widespread in human reasoning and communication (see Brandom 2008; Pollock 1987; Horty 2012; Veltman 1996).

Despite this pervasiveness, many features of nonmonotonic reasoning remain to be explored. In this paper I focus on its occasion-sensitivity – which, I think, is not yet well understood.

By saying that defeasible inferences are occasion-sensitive I mean that their *goodness is sensitive to specific features of each episode of inferring* – in particular, features of the circumstances surrounding the *situation targeted by the inference* (i.e. the situation the inference is about), but *possibly* also features of the *context in which the inference is made or assessed*.<sup>1</sup> I will say that an account of defeasible inferences is *particularist* if it introduces this sort of occasion-sensitivity *in a way that cannot be easily captured by general rules*.

In a nutshell, my proposal is that the distinction between potential defeaters that have to be discarded by the premises of an inference and potential defeaters that may remain unconsidered (without making the inference enthymematic) is occasion-sensitive. Typically, there are many possible defeaters for the goodness of nonmonotonic inferences; however, in general such inferences can remain non-enthymematic without the introduction of additional premises ruling out each of these defeating conditions. I will argue that only those defeaters whose possibility is *relevant* will need to be considered explicitly. Whether the possibility of a given defeater is relevant, I will claim, depends on the circumstances surrounding the particular situation the inference is about, and plausibly also on features of the context in which the inference is made or assessed. A defeating possibility that can remain unconsidered in certain circumstances of inference, may need

---

<sup>1</sup> I reserve the notion of context to refer to the circumstances in which inferences are either *made or assessed*, rather than to the circumstances surrounding the *situation the inference is about*. My reason for this usage is that it reflects how the notion of context-sensitivity is typically understood in debates in semantics. So, for example, subject-sensitive invariantism about knowledge attributions (Hawthorne 2004; Stanley 2005) is not taken to be a contextualist theory, even if it makes the accuracy of knowledge attributions dependent on the circumstances of the subject *targeted* by the attribution.

to be explicitly discarded by the premises of the inference in some other circumstances. Thus, if successful, my analysis will show that the same (nonmonotonic) inferential entailment may be good on some occasions and bad or enthymematic on others – in which further premises ruling out certain possible defeaters need to be added. In other words, I will argue that it is an occasion-sensitive matter whether a given inference is good all things equal (i.e. absent defeaters), or whether it is rather an enthymeme.

I will present, therefore, an occasion-sensitive characterization of defeasible inferential entailments. My main goal is to examine the consequences of this occasion-sensitivity for inferentialist theories of content (for instance, Brandom 1994; Harman 1982). According to inferentialism, the content of a proposition is determined by its inferential connections with other propositions. Most inferentialist theories include nonmonotonic inferences among those inferences that determine semantic content (Brandom 2000). Thus, if nonmonotonic inferential entailments are occasion-sensitive – as I will argue – then inferentialist theories of content involving nonmonotonic inferences will become occasion-sensitive as well.

It is not my purpose in this paper to offer arguments in favor of inferentialist theories, but rather to examine the occasion-sensitive dimension that is introduced in these theories when defeasible inferences are in play. If my analysis is accurate, the relations between inferentialism and particularist perspectives on human language and thought (for instance Dancy 2004; Travis 2008) will turn out to be deeper than what has been usually acknowledged: on the view I will explore, (nonmonotonic) inferential entailments among propositions cannot be specified in abstraction from the circumstances surrounding each particular occasion of inference. At any rate, the occasion-sensitivity of defeasible reasoning revealed by this discussion is, I think, interesting on its own – beyond its consequences for inferentialism – and may contribute to illuminating current debates in epistemology and metaethics.

I start by offering a brief overview of inferentialist theories of content, emphasizing those

aspects relevant for the ensuing discussion. In section 3, I present a first form of contextual dependence affecting inferentialist theories – due to the influence of collateral premises. After that, in section 4, I introduce nonmonotonicity. Next, in section 5, I explore the occasion-sensitivity of nonmonotonic inferential connections. In section 6 I discuss the consequences of this occasion-sensitivity for inferentialist theories.

## **2. Inferentialist Theories of Content**

According to inferentialist theories, the content of a proposition is determined – at least partly – by its inferential connections with other propositions – that is, what other propositions entail it and what further propositions are entailed by it. For instance, the content of the claim "The table is red" would be partly determined by its entailing the claim "The table is colored". For the purposes of this paper, I do not have to decide whether propositional content is constituted or merely determined by inferential relations (so I leave open the possibility that, even if content depends on the inferential role of propositions, it is not identified with it).

There are many versions of inferentialist theories of content (for instance, Brandom 1994; Harman 1982; Sellars 1953), but I will not need to consider most of their differences. I will only assume, following Brandom (1994, 2000) – who in turn follows Sellars (1953) –, that the inferential connections a proposition is involved in are directly related to the role played by such a proposition in reasoning and, more generally, in argumentative practices of giving and asking for reasons. On this view, an inferential link among propositions holds if the piece of reasoning underlain by it is good – that is, if someone having good reasons to accept the premises has, by virtue of this, good reasons to accept the conclusion. So, if  $p$  entails  $q$ , then having good reasons to accept that  $p$  (e.g. to assert it or believe it) gives one good reasons to accept that  $q$ .

In order to work with a specific inferentialist proposal, I will take Brandom's (1994) theory as my model in this paper, in part because it is one of the most worked out renderings of

inferentialism and in part because it offers the most suitable framework for the discussion of the sorts of considerations I will tackle— in particular, Brandom explicitly discusses defeasible reasoning and acknowledges the nonmonotonicity of most inferential connections (Brandom 2000; 2008). However, I will not need to follow all the details of Brandom's theory and I will not claim to be offering a complete and accurate representation of his views. At any rate, the conclusions I will reach apply as well to other sorts of inferentialism.

In the next section, I present a first approach to a contextualist analysis of inferentialism. More specifically, I discuss how inferentialism acquires a holistic dimension as a result of the dependence of the inferential consequences of a proposition on the collateral premises available; I will argue, however, that the sort of context-dependence associated with this holism is rather limited and constrained. In the rest of the paper, I try to show that when nonmonotonic inferential connections are allowed, a more radical kind of occasion-sensitivity arises.

### **3. Inferential Holism**

In general, what follows from a given proposition depends on the collateral premises available. What is entailed by  $p$  may differ from what is entailed by  $p$  and  $q$ . So, the proposition "The solution to the equation is a natural number between 4 and 7" only entails the conclusion "The solution to the equation is 6" if it is joined by some suitable collateral premise – for instance, "The solution to the equation is a multiple of 3." In the same way, the claim "The clothes will get wet" follows from the premises "The clothes are drying outside" and "It will rain", but not from the first premise alone.

As a result of this well-known fact, the inferential entailments a given proposition is involved in depend on what other propositions can be used as collateral premises. Let me define – in accordance with Brandom (1994: 480-81) – the *inferential significance* of a proposition  $p$  as what further propositions follow from  $p$  (maybe in conjunction with other premises), and from what other propositions  $p$  follows. The idea I am presenting, therefore, is that the inferential significance of a

proposition depends on what other propositions are available as collateral premises – what I will call the *inferential setting*. The introduction of a new collateral premise would change the inferential setting and, as a result, would potentially modify the inferential significance of every other proposition.

Thus, inferentialism is a holistic theory, in the sense that the inferential significance of a proposition depends, at least in principle, on its relation to every other proposition available as a premise (see Brandom 1994: 477-90; for criticism, Fodor and Lepore 2001, 2007). In an inferentialist framework in which the content of beliefs is taken to be given by their inferential significance, this would have the consequence that anytime an agent becomes entitled to accept a new belief her other beliefs start expressing a different content. This may seem implausible – in particular, it would mean that agents with different collateral beliefs cannot share a belief with the same content (see Fodor and Lepore 2001, 2007).

I do not think, however, that this sensitivity to the inferential setting actually carries such radical consequences. In particular, it is still possible to have a notion of content which is constant across inferential settings – that is, which does not vary when new collateral premises are added. In order to do so, I will follow a suggestion made by Brandom (1994: 482-84) and characterize the content of a proposition as a function from an inferential setting (i.e., a set of other propositions available as collateral premises) to an inferential significance. Thus, grasping the content of a proposition would amount to knowing what follows from it – its inferential significance – given the collateral premises available. Accordingly, the content expressed by some belief will not depend on what other beliefs are accepted by the agent.

In the next sections I will argue that a further sort of holism and occasion-sensitivity – whose consequences remain far less familiar – stems not from the fact that new premises may lead to new conclusions, but rather from the fact that the introduction of new premises may defeat one's entitlement to conclusions that were previously accessible. Inferential nonmonotonicity is the

source of an interesting kind of semantic occasion-sensitivity. I start by exploring some features of nonmonotonic inferences.

#### 4. Nonmonotonic Inferences

Monotonic inferences are those inferences whose goodness cannot be defeated by the introduction of new premises (consistent with the old ones): after introducing a new collateral premise, one can still infer everything that could be inferred before. It is customary to take formal validity according to classical (monotonic) logic as the paradigmatic model of inferential goodness. As a result, there is the temptation to think that valid inferences are typically monotonic. Think of a basic example:

$$\begin{array}{l} (1) \quad p \\ \quad \quad q \\ \quad \quad \text{-----} \\ \quad \quad p \ \& \ q \end{array}$$

(1) will remain good (in classical logic) no matter what extra premise “*r*” is added.

However, things change when one takes into consideration *materially good inferences*, that is, inferences that are good in virtue of the content of the propositions involved – rather than their formal structure. Clearly, an inferentialist theory of content has to appeal to this sort of inference. Formally valid inferences will not help to characterize the content of non-logical concepts, since such inferences are valid because of their formal structure, with independence of the particular contents of the non-logical concepts involved in them (Sellars 1953; Brandom 1994: 97-105).

Some material inferences seem to be monotonic – even if they are not valid in virtue of their form. Think, for instance, of the inference from “France is bigger than Monaco” to “Monaco is smaller than France”, or the inference from “Tim arrived before Tom” to “Tom arrived after Tim.” Other material inferences, on the contrary, are clearly nonmonotonic – as Brandom points out (see

Brandom 2000: 87-89).<sup>2</sup> So, the inference represented in (2) is typically regarded as good:

(2) It is raining

---

The street will get wet

However, the introduction of the new premise “The street is covered by a canopy” would turn (2) into a bad inference. After the introduction of that additional premise, the inference that becomes good is (3):

(3) It is raining

The street is covered by a canopy

---

The street will remain *dry*

Another example is:

(4) The match is struck

---

The match will light

(4) will become bad if one adds the premise “The match is wet” or “There is no oxygen in the room.” When an inference becomes bad due to the addition of a new premise, this premise is said to be *defeating* the goodness of the inference. Accordingly, such premises are usually called *defeaters*

---

<sup>2</sup> Nonmonotonic inferences have been studied in logic and Artificial Intelligence (Reiter 1980; McCarthy 1986; Veltman 1996) and also in connection with defeasible reasoning in epistemology (Pollock 1987; Williams 2001) and in metaethics (Horty 2012; Thomas 2012; Schroeder 2011; Vayrynen 2004).

(see Horty 2012; Dancy 2004; Pollock 1987).<sup>3</sup> In particular, I will focus on undercutting or disabling defeat – that is, cases where, given the presence of the defeater, the premises stop counting as reasons in favor of accepting the conclusion.

Far from being an anomaly, nonmonotonic inferences seem to be the standard case in human practices, at least when one moves outside the special domains of mathematics, maybe fundamental physics and, of course, monotonic logic. Most inferences in our daily lives, and also in scientific reasoning turn out to be nonmonotonic (Brandom 2008; Pollock 1987). Monotonicity has tended to be seen as a default feature of good inferences only because of a formalist bias in the main tradition that has studied inferential validity. Inferentialist semantics, however, usually include nonmonotonic inferences among the inferential connections relevant for the determination of the content of propositions. In this paper, I focus on this sort of version of inferentialism: I will take for granted that, at least for a large class of propositions, the inferences characterizing propositional content include nonmonotonic ones.

It may be argued that nonmonotonic inferences are actually enthymemes, which can be made formally valid by adding a suitable conditional premise. So, (2) would remain enthymematic unless one adds the conditional premise “If it is raining, then the street will get wet”. Sellars complained against this move as a 'received dogma' (1953: 313); in a similar vein, Brandom offers a pragmatist motivation for rejecting this strategy: the formalist reaction merely 'trades primitive goodness of inference for the truth of conditionals' and, according to him, inferentialism is appealing to the pragmatist precisely because the grasp of inferential goodness is easier to explain as a practical skill than the grasp of truths (1994: 98). It can be added that speakers generally tend to accept (good)

---

<sup>3</sup> This form of defeat is usually contrasted with rebutting defeat, where the original premises still offer reasons in favor of the initial conclusion, but these reasons are outweighed by those provided by the defeater in favor of alternative conclusions (see Pollock 1987; Schroeder 2011; Dancy 2004).

material inferences as correct. At any rate, this formalist retreat is not available to the inferentialist, since the resulting inference would be good in virtue of its logical form and, therefore, would not be of use in specifying the contents of non-logical concepts. As I am assuming an inferentialist framework, I will take for granted that not every material inference is enthymematic.

Another suggestion is that one may restore monotonicity by adding auxiliary anti-defeater premises to a nonmonotonic inference. These auxiliary premises would protect the inference from possible defeaters, without making the resulting inference logically valid.<sup>4</sup> One way in which this can be done is by adding a premise that is incompatible with the presence of the defeater. In this way, the auxiliary premise “The match is dry” can be added to the premise “The match is struck” in order to secure the conclusion “The match will light” (defeaters have to be consistent with the original premises, otherwise they just make the inference unsound). A second way in which an auxiliary premise may protect an inference from a potential defeater is by undermining its defeating power. For instance, “The canopy is full of holes” disables the defeating power of “The street is covered by a canopy” with respect to (2). I will say that a defeater is itself defeated when the addition of a further new premise allows one to infer again the original conclusion (this is what happens if the premise “The canopy is full holes” is added in (3)).

The strategy of adding protective premises will surely work when facing any given specific defeater. But, according to Brandom (2000, 2008), the problem is that, in many cases, there are an indefinitely large number of possible defeaters. So, even if the premise “The match is dry” is added, the match may still fail to light due to the absence of oxygen in the room, to an extremely low temperature or due to some other unforeseen obstacle. In principle, there is no way of determining when the list of possible defeaters will end: with some imagination, there is always the possibility of finding further defeating conditions (some evil magician may be casting a spell on the match).

---

<sup>4</sup> Anti-defeater premises are what Dancy (2004) calls ‘enablers’, that is, conditions that allow other premises to count as reasons for some conclusion.

I will assume that Brandom is right in holding that the list of defeaters is, at least in many relevant cases, indefinitely large. Nonmonotonicity, in these cases, would be irreparable. Should one say that these intrinsically nonmonotonic inferences are doomed to remain enthymemes (i.e. there will always be some unstated anti-defeater premise)? I think that the most reasonable answer is negative. Inferences of this sort are usually taken as good in our practices, even if there is no hope – under the present assumption – of bringing to light every unstated anti-defeater. So I think that the best option is to accept the nonmonotonic nature of this sort of inference and consider them good as they stand – albeit *defeasibly* good. After all, I would say that the notion of enthymeme is only useful when the relevant hidden premises may be uncovered, so that at some point the inference stops being enthymematic.

The picture that emerges from this discussion is that good material nonmonotonic inferences are acceptable *even if not all anti-defeater considerations figure among the collateral premises* (remember that I am assuming that it is not possible to state *all* such considerations). In this way, the goodness of the inference will be underpinned by an indefinitely large number of anti-defeater considerations that remain unstated in the background – for instance, the condition that the temperature is not extremely low, that there is an atmosphere with oxygen or that no evil sorcerer is casting defeating spells. When one of these conditions fails to obtain, a defeater is introduced and the inference becomes bad (unless the defeater is itself defeated by some further fact).

According to this picture of defeasible reasoning, agents do not need to provide positive reasons for discarding the presence of all possible defeaters, in order to count as entitled to endorse the conclusions of some inference. In the next section I discuss what determines whether a given defeating possibility needs to be explicitly tackled by the premises of a good, non-enthymematic inference. My contention will be that this is an occasion-sensitive matter.

## **5. The Occasion-Sensitivity of Inferential Goodness**

Which defeating possibilities may remain unconsidered in the background, and which have to be explicitly considered by the premises, on pain of making the inference enthymematic? A quick answer is that only defeating possibilities that are *relevant enough* on the occasion for inferring need to be considered, whereas less relevant possibilities may be properly ignored. I think that this answer is essentially right, but it will remain rather uninformative unless something more substantial is said about what makes a defeating possibility relevant enough on a given occasion of inference.

At this point, one could turn to the sorts of rules proposed by Lewis for properly ignoring possibilities (Lewis 1996; for discussion, see Blome-Tillmann 2009; McKenna 2014). Although Lewis' rules concern knowledge attributions, I think it is plausible that similar considerations apply to assessments of inferential goodness; after all, inference is one of the standard ways of acquiring knowledge.<sup>5</sup>

Regardless of the exact relation to Lewis' rules, my own proposal is that a defeating possibility may be properly ignored by the premises of an inference only if there are *no reasons for the agent to suspect that the defeating possibility may obtain* (or, less stringently, only if such reasons are weak enough).<sup>6</sup> In situations where there actually exists some such reason, the defeating

---

<sup>5</sup> Ichikawa (2011) argues that knowledge attributions and counterfactuals rely on the same rules for ignoring possibilities (see also Lewis 2016). Plausibly, there is a close relation between the possibilities relevant for assessing counterfactual inferences and the possibilities relevant for assessing counterfactual conditionals.

<sup>6</sup> A possible view is that defeasible inferences are (non-enthymematically) good when conditions are *normal*. In order to flesh out this view, an account of normality would be needed. My proposal can be seen as a characterization of the relevant form of normality in terms of absence of reasons to suspect the presence of defeaters (for discussion of the notion of normality in relation to defeasible reasoning, see Cullity 2002; Nickel 2010).

possibility becomes relevant, and it will become necessary to introduce a corresponding anti-defeater consideration in the premises (I will treat doubting and suspecting as duals: if there are reasons to suspect that some defeater obtains, then there are reasons to doubt that it is absent). For instance, if there are reasons to doubt that there is oxygen in the surroundings, the agent may be required to show that this is so before being licensed to conclude that the match will light when struck. But, as long as the agent does not possess reasons to suspect that some defeater may obtain, she will be entitled to treat her inference as non-enthymematically good (although this is a defeasible entitlement, which will be revoked should such reasons appear).

I will not try to give an exhaustive characterization of the sources of reasons to suspect that a defeater may be present. I will just examine certain circumstances in which it is plausible to think that a defeating possibility is relevant. More specifically, I submit that there are reasons for suspecting the presence of a given defeater (which is thereby relevant) in circumstances where such a presence is likely enough or rationally expectable.<sup>7</sup> In the following sections, I develop this

---

<sup>7</sup> An interesting question is how propositions *about defeaters* should be dealt with on an inferentialist approach (e.g. the proposition “This defeater is relevant enough”). A natural option is to think of terms like ‘defeater’ as playing an expressive role: such vocabulary would allow us to talk and argue explicitly about the inferential structure of our discursive practices (in particular, about the conditions in which some inference is good). This story would be in line with Brandom’s views about the expressive role of logical, modal or epistemic vocabulary (see Brandom 2008). The content of propositions asserted using this expressive vocabulary would be characterized in the standard inferentialist way – that is, in terms of the inferences such propositions are involved in (e.g. the inference from “This defeating possibility is relevant enough” to “This defeating possibility cannot be properly ignored” or the inferences associated with the analogous of Lewis’ rules). I thank an anonymous reviewer for drawing my attention to this issue.

suggestion in more detail, showing how it makes inferential goodness depend on surrounding features of the occasion of inference.

### 5.1. Typical conditions

Arguably, when the presence of a certain defeater is typical in the sort of situation the inference is about, the possibility of such a defeater is relevant and has to be explicitly considered. More precisely, I will treat a possible defeater as relevant when there is some fact (maybe ignored by the agent) such that, in light of that fact, the presence of the defeater would be likely from the agent's epistemic perspective (i.e. given what she knows): *if added to the rest of the agent's evidence*, such a fact would give her reasons for suspecting that the defeater may obtain.

In this way, when the defeater does occur, it is always relevant, since (trivially) *if the fact that the defeater actually obtains is added to the agent's epistemic perspective, its presence becomes maximally likely from that perspective.*<sup>8</sup> Of course, the reasoning agent may remain unaware of the presence of the defeater. This only means that, even if there is a very good reason to think that the defeater obtains, such a reason may be beyond the agent's ken (it may not be one of her possessed reasons). Still, since the defeating possibility is actually relevant, an inference ignoring such a defeating possibility will not be good. This, I think, is a welcome result, since it blocks the possibility that an inference from true (but undercut) premises to a false conclusion counts as good, just because the agent ignores the presence of the defeater.

Likewise, the fact that a defeating possibility is *environmentally typical* makes that possibility relevant, and thereby in need of being ruled out or itself defeated by the premises of the inference, even in cases where the agent is not aware of such typicality. If the fact that a given defeater is environmentally typical were included in the agent's evidence, she would possess reasons for suspecting the presence of that defeater – at least unless she had further reasons to

---

<sup>8</sup> This coincides with Lewis' rule of actuality (1996: 554).

believe that, despite this typicality, the defeater is absent in the specific situation targeted by the inference.

Consider the following inference:

(5) The water is boiling

---

The water is very hot (over 70°C)

Given the standard atmospheric conditions of the Earth's surface, this inference is good. However, it is nonmonotonic: if the atmospheric pressure dropped dramatically, the inference would be defeated. The introduction of the additional premise "The pressure is extremely low" would turn it into a bad inference (the boiling temperature of water drops as pressure decreases). This does not mean, nevertheless, that I always have to add a protective anti-defeater premise stating that pressure is not extremely low, in order for the inference to be good. In ordinary cases, being entitled to claim that the water is boiling is enough to count as having good reasons to conclude that the water is very hot.

By contrast, when the pressure happens to be atypically low, this fact needs to be included as a premise, if I want to be entitled to conclude that, since the water is boiling, it is *not* very hot. In an ordinary Earthling environment, in an atypical situation where the pressure is low, inference (6) is enthymematic:

(6) The water is boiling

---

The water is *not* very hot (below 50°C)

Now, imagine that a low pressure laboratory where water is typically kept at very low pressures, so that it boils at mild temperatures (alternatively, imagine people living on the surface of Mars). Plausibly, in a linguistic community living in this environment, inference (6) would be taken as (defeasibly) good. Note that, for them, very hot boiling water is as unexpected as cold boiling water is for us.

Of course, if the pressure goes unusually high, up to Earth's levels, the inference will be defeated. But this possibility would not make it the case that, in general, one should always add the collateral premise "The pressure is very low". As a general rule, in this environment, that condition may remain in the background, as an unstated anti-defeater condition – in the same way that, in my Earth's kitchen context, the condition that the pressure was not very low could remain in the background. So, the inferential link (6) is good, in this environment, without needing to be complemented by the collateral premise "The pressure is very low". On the contrary, in the Earth's surface scenario, (6) was bad unless joined by such a collateral premise – *even in unusual situations in which the pressure happened to be as low as in the second scenario*.

What about Jane, who is not aware that she is visiting a low pressure laboratory? As she ignores that she is in such an unusual environment, the fact that her ordinary experience with boiling water has been in high pressure conditions gives her reasons to expect that the pressure will be high: defeating high pressures are likely given her epistemic perspective. Therefore, (6) will be enthymematic for her – she would need to add a premise stating that the pressure is extremely low.

Finally, imagine an exceptional situation in which one of the pots with boiling water in the laboratory is at high pressure. Even in this unusual situation, where the boiling water actually happens to be hot, inference (5) will be seen as bad – as enthymematic – from the point of view of this second environment. In order to vindicate the goodness of (5), it is necessary to introduce the collateral premise "The pressure is around 1 atmosphere".

Arguably, this last case gives rise to inferential counterparts of fake barn scenarios. Think

again of Jane. Imagine that, out of luck, she happens to see the only pot boiling at a high pressure in the laboratory. Still ignoring that she is in a special laboratory, she goes on to infer (as she would in an ordinary context, like her kitchen) that the water in the pot is very hot. Plausibly, this is not a good inference. The possibility of low pressures is relevant enough and has to be explicitly ruled out, given the fact that she is in the special laboratory *and that she does not know that the specific pot she is facing is exceptional*, in that it is the only pot in the facilities at high pressures. This is analogous to what happens in fake barn cases, where the agent is taken not to know that she is facing a barn despite being in front of one of the only real barns in the county, because she is not able to distinguish it from the fakes in her surroundings.<sup>9</sup>

The upshot of the boiling water example is that the same inferential link can be good in some environment but not in another, even if the specific situation targeted by the inference is the same in both cases (that is, even if in both cases the inference is about a pot with boiling water at the same pressure and temperature). In this way, (5) is (defeasibly) good in an ordinary Earth's surface scenario, but bad in a low pressure environment such as a special laboratory or Mars' surface – even in an exceptional situation where pressure happens to be high there. Conversely, (6) is enthymematic in ordinary Earthling contexts but (defeasibly) good in environments where low pressures are typical. So, the goodness of an inferential connection is not only determined by what propositions are linked, but also by the surroundings of the [situation targeted by the inference](#) – in particular, by whether defeating possibilities are typical in such surroundings.

Nonmonotonic inferential goodness, thus, depends holistically on [aspects](#) of the

---

<sup>9</sup> At least, this is the mainstream reading of fake barn cases. Nevertheless, some authors claim that there is empirical evidence that people tend to attribute knowledge to subjects in fake barn examples (Colaço et al. 2014; Turri et al. 2015; for disagreeing empirical results, see Nagel et al. 2013). This empirical evidence, however, focuses on cases of perceptual knowledge. It would be interesting to see whether people have similar intuitions in the inferential cases I am considering.

circumstances of inference. The point is not merely that inferential connections can be defeated, but rather that which possible defeaters need to be explicitly considered by the premises of an inference (on pain of making it enthymematic) depends on the circumstances surrounding the situation the inference is about. So, whether some conclusion (defeasibly) follows from some premises cannot be determined merely by considering a general inferential schema, without being sensitive to specific features of the particular episode of inferring. In this sense, what (defeasible) inferential connections between propositions hold is an occasion-dependent issue (remember that I am saying that inferential goodness is occasion-sensitive if it depends either on features of the context in which the inference is made (or assessed) or on the surrounding circumstances of the particular situation targeted in that episode of inferring).

So far I have argued that the relevance of defeating possibilities depends on features of the situation the inference is about. Plausibly, whether some defeating possibility is relevant enough will also depend on features of the *context in which the inference is made* (or perhaps assessed): for instance, the sort of information available, the history and purposes of the conversation or the kind of inquiry the relevant agents are engaged in. Depending on these contextual factors, different defeaters may count as relevant enough. I turn now to explore such context-sensitivity.

### 5.2. *The context-sensitivity of inferential relations*

I have claimed that a defeating possibility becomes relevant (and needs to be considered by the premises of the inference) when there are (sufficiently strong) reasons to suspect that some defeating condition could obtain. But, how strong do these reasons to suspect need to be in order to make a defeating possibility relevant *enough*? What is the threshold of relevance for a defeating possibility to merit consideration?

An interesting possibility is that the answer to these questions will be sensitive to features of the *context in which the inference is made*, for instance the purposes of the conversation and the

type of inquiry agents are engaged in (see Williams 2001); in particular, it is plausible to think that the threshold of relevance will depend on pragmatic factors such as the costs of endorsing wrong conclusions. In [contexts](#) in which the practical costs associated with inferential errors are high, one may be required to exclude explicitly defeating possibilities that in more relaxed contexts could remain unconsidered.

This is related to the (increasingly popular) view that the practical ‘encroaches’ on the epistemic (see, among others, Fantl and McGrath 2009; Stanley 2005; Hawthorne 2004;). According to this view, the epistemic status of an agent depends on pragmatic factors, such as how high stakes are. In particular, pragmatic factors would have a say in determining how much evidence (ruling out how many alternative possibilities) an agent needs to possess in order to count as knowing some proposition. In high stakes situations, more evidence would be required than in low stakes contexts.

One of the standard ways of acquiring knowledge is through evidence-based reasoning, which, in general, involves defeasible inferences. Therefore, if the goodness of defeasible inferences depends on pragmatic factors, this will explain the sensitivity of the corresponding knowledge attributions to such pragmatic considerations. In this way, a suitable analysis of defeasible reasoning (to which I am trying to contribute here) will allow for a clearer and more systematic appraisal of the epistemological debate on pragmatic encroachment (a detailed account of the relations between pragmatic encroachment and defeasibility, however, lies beyond the scope of this paper).

Arguably, contextual factors may determine not only the threshold of relevance required for a defeating possibility to deserve consideration, but also *how relevant* the possibility of a given defeater is. Lewis (1996) has contended that merely mentioning some possibility (or otherwise paying attention to it) is enough to make it relevant in the context of the conversation. Without going that far, one could think that features of the conversational context have a role in determining

the strength of the reasons to suspect that a certain defeater may obtain.<sup>10</sup> In particular, it seems that the information available to agents can affect whether some defeating possibility is rationally expectable – for instance, it can reduce the relevance of some otherwise highly relevant possibility.

Consider the following case. Mary lives in a world where jellyfish are usually venomous – most jellyfish she has encountered produce a painful sting when touched. Imagine now that she happens to find herself in a peculiar island where jellyfish are generally harmless (venomous jellyfish are extremely rare there). Nevertheless, if Mary is unaware of the peculiarities of the island she is in, she will still have reasons to suspect that the jellyfish she has just accidentally touched while swimming may be venomous. This is a relevant, likely possibility, given that jellyfish in most parts of Mary's world are venomous, and that she does not know that she is in a peculiar place. By contrast, for Tom, a native islander who is familiar with the peculiarities of the island's fauna, the possibility that the jellyfish in the beach are venomous may be ignored (it is an unlikely possibility in light of Tom's experience). For instance, such a possibility would not need to be considered in order for Tom to be entitled to conclude that it is safe to swim in the beach since only jellyfish (but not sharks) are to be encountered.

In fact, it might be argued that some defeating possibility can be made relevant even by *misleading* information, provided that such information comes from a generally reliable source and is convincing. Given the fact that the information is coming from usually reliable sources, the presence of the defeater is rationally expectable for the agent (even if the defeating condition was not actually likely to obtain, since the information was misleading).

### 5.3. *Summary: occasion-sensitive inferential entailment*

---

<sup>10</sup> I will not commit myself either to Lewis' rule of attention (1996, p. 559). The fact that we pay attention to a certain possibility does not necessarily mean that there are good reasons to think that the possibility may obtain (see McKenna 2014).

Although there is room for debate about how to interpret the details of the examples I have given, I think that the discussion above makes a plausible case for the view that the goodness of inferential connections is sensitive to the occasion of inference – both to environmental features of the situation targeted by the inference and to features of the context in which the inference is made. Inferential goodness, therefore, cannot be determined ‘in vacuo’, just by considering the propositions serving as premises and conclusion; in order to be able to ascertain what is a good reason for what, what follows (defeasibly) from what, one needs to be sufficiently responsive to [different aspects of the circumstances surrounding the specific episode of inferring](#). Otherwise, it will not be possible to tell whether a given inference is good or rather enthymematic. More specifically, [one needs to be able to appreciate when some defeating possibility is relevant enough](#). [Arguably, this ability is not merely a matter of grasping general rules, but rather involves attending to occasion-specific features of the circumstances of inference \(such as the typical environmental conditions or the conversational stakes\); thus, we get a particularist picture of inferential goodness.](#)

These conclusions can also be framed in terms of default reasoning. Defeasible inferences are sometimes characterized as default reasoning rules, that is, rules of reasoning that can be applied if they do not conflict with stronger rules and are not disabled by some defeating consideration (Horty 2012). An example would be the default codified by the conditional “If the water is boiling, it is very hot” or by the generalization “Birds fly”. The view I am putting forward is that it is an occasion-sensitive matter whether some reasoning rule is to be treated as a default, or whether it should be taken to be enthymematic and in need of complementation by further premises. A given reasoning rule cannot be identified as a default, rather than as an enthymeme, independently of the circumstances in which that rule is to be applied (see McKeever and Ridge 2007). For instance, “If Perry is a mammal, then Perry does not lay eggs” arguably instantiates a default rule in most ordinary contexts, but it would not count as such in an environment where platypuses are very common (one would need the extra premise “Perry is a platypus”). So, features of the circumstances

of application – including the environmental typical conditions and pragmatic factors such as how high stakes are – would determine not merely whether some default rule is disabled or outweighed, but also whether it counts as a prima facie default to start off.

I will discuss now how inferentialist theories of content are affected by the occasion-sensitivity of inferential connections.

## **6. Consequences for Inferentialism**

Inferentialist theories of the sort I am interested in maintain that the content of propositions is determined by the inferential relations such propositions are involved in, which are often nonmonotonic. In turn, I have argued that whether a (nonmonotonic) inferential relation among propositions holds is an occasion-sensitive matter. Therefore, nonmonotonicity turns inferentialism into a particularist, occasion-sensitive theory of content. Note that I am not referring to the fact that the inferential significance of a proposition depends on the collateral premises available – such dependence is acknowledged by Brandom (1994: 477-90) and was discussed in section 3. The sort of occasion-sensitivity I am exploring here goes farther: given a set of available collateral premises, what follows from a certain proposition depends on features of the surrounding circumstances of inference, because such features determine whether some defeating possibility may remain ignored or needs to be explicitly considered by the premises.

One may wonder whether this particularist framework allows for the identification of a workable notion of propositional content, which makes communication across contexts possible. A first option is to think of inferentialism as a generalized version of what MacFarlane (2014) calls *indexical contextualism*. On indexicalist models, utterances of the same sentence express different propositions in different contexts. This is what happens typically with sentences involving indexicals or demonstratives. What proposition I manage to express by asserting ‘This street is closed to traffic today’ depends on the day of utterance and on what street is picked out by the

demonstrative 'this'.

If inferentialism is constructed as a form of indexicalism, any assertion (involving concepts caught up in nonmonotonic inferences) will express potentially different contents on different occasions of utterance. For instance, the term 'water' would express different concepts when talking about water in my kitchen and when talking about low pressure laboratories – and maybe also when talking in high and low stakes contexts.

This consequence seems undesirable if one thinks that the possibility of agreement and disagreement – and, more broadly, the possibility of communication – requires that speakers can share and transmit contents across different occasions of utterance. But, perhaps, one could think of communication as involving a form of coordination among agents that does not rest on sameness of content (indeed, Brandom seems to gesture towards such an alternative conception of communication, see 1994: 485-90).

Anyway, if it is desired to preserve the possibility of shared contents across contexts, one can always iterate the strategy rehearsed in section 3, and treat propositional content as a function – this time, a function with an inferential significance as an output and taking as inputs the inferential setting (i.e., collateral premises available) *and a defeater-relevance parameter* (keeping track of what defeating possibilities are relevant enough and cannot remain unconsidered in the background). Thus, in order to assess the inferential significance of some proposition, in addition to knowing which collateral premises are available, one would also need to be sensitive to those cotextual factors that determine what defeating possibilities may be properly ignored; then, it would be possible to ascertain what follows from that proposition and what entails it. The content expressed by some assertion would not change with changes of the relevance parameter, but the inferential significance attached to that content would.

This second approach can be seen as an inferentialist form of *relativism*. On relativist views, the truth value of a given proposition is relativized to *parameters in the circumstances of*

*evaluation*, even if the proposition itself remains the same for different values of these parameters (see MacFarlane 2014; Kölbel 2004). In standard semantics, the circumstances of evaluation are typically identified with the world (and maybe time) of evaluation. But, according to relativists, additional parameters could in principle be included in the circumstances of evaluation. For instance, one could add some epistemic, aesthetic or gastronomic standard. In this way, the truth of propositions about the tastiness of apples would change when evaluated in relation to the standards of taste of different individuals.

In truth-conditional frameworks, the truth value of a proposition determines the correctness of its assertion; analogously, on inferentialist views the inferential significance of a proposition determines the correctness of the inferential moves involving it. So, in the same way that in relativistic truth-conditional theories the truth value of propositions is relativized to contextual parameters, in inferentialist versions of relativism it would be the inferential significance of propositions that is relativized to contextual factors. In particular, in the sort of relativism I am considering here, the inferential significance of a proposition would be made to depend on those occasion-sensitive factors that determine what defeating possibilities may be properly ignored.

My aim here has not been to adjudicate between indexicalist and relativist characterizations of inferentialist propositions – this is a choice that will depend on one's views about communication and about the role played in it by the notion of content. I have just tried to present two possible ways of modelling the kind of inferentialist contents resulting from the discussion in this paper. Note, anyway, that both in indexicalist and relativist models the use of inferentialist contents in one's language and thought would be underlain by occasion-sensitive skills. On indexicalist views such skills would be needed to see what (invariantist) propositions are expressed on each occasion of use of a sentence; meanwhile, relativist approaches would rely on occasion-sensitive skills in order to ascertain the inferential significance of a given proposition in different circumstances.

## 7. Conclusions

My main thesis has been that nonmonotonicity introduces a far-reaching occasion-sensitive dimension in inferentialist theories of content – at least in those versions of inferentialism in which the content of propositions is determined by their involvement in nonmonotonic inferences. For those who think that adequate theories of content should not exhibit such occasion-sensitivity, this result will count against the plausibility of inferentialism; others, by contrast, may take it to reveal an interesting feature of human language and conceptual thought.

## References

- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118, 241-294.
- Brandom, R. (1994). *Making it Explicit*. Cambridge, MA: Harvard University Press.
- Brandom, R. (2000). *Articulating Reasons*. Cambridge, MA: Harvard University Press.
- Brandom, R. (2008). *Between Saying and Doing: Towards an Analytic Pragmatism*. Oxford: Oxford University Press.
- Colaço, D., W. Buckwalter, S. Stich & E. Machery (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(02), 199-212.
- Cullity, G. (2002). Particularism and Presumptive Reasons. *Aristotelian Society Supplementary Volume*, 76, 169-190.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Clarendon Press.
- Fantl, J. & M. McGrath (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Fodor, J. & E. Lepore (2001). Brandom's Burdens. *Philosophy and Phenomenological Research*, 63, 465-481.
- Fodor, J. & E. Lepore (2007). Brandom Beleaguered. *Philosophy and Phenomenological Research*, 74, 677-691.

- Harman, G. (1982). Conceptual Role Semantics. *Notre Dame Journal of Formal Logic*, 23, 242-256.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Horty, J. (2012). *Reasons as defaults*. Oxford: Oxford University Press.
- Ichikawa, J. (2011). Quantifiers, knowledge, and counterfactuals. *Philosophy and Phenomenological Research*, 82(2), 287-313.
- Kölbel, M. (2004). Indexical relativism versus genuine relativism. *International journal of philosophical studies*, 12, 297-313.
- Lewis, D. (1996). Elusive knowledge. *Australasian journal of Philosophy*, 74, 549-567.
- Lewis, K.S. (2016). Elusive Counterfactuals. *Noûs*, 50(2), 286-313.
- McCarthy, J. (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28, 89-116.
- MacFarlane, J. (2014). *Assessment sensitivity: Relative truth and its applications*. Oxford: Oxford University Press.
- McKenna, R. (2014). Normative scorekeeping. *Synthese*, 191, 607-625.
- McKeever, S. & Ridge, M. (2007). Turning on Default Reasons. *Journal of Moral Philosophy*, 4(1), 55-76.
- Nagel, J., V. San Juan & R.A. Mar (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129(3), 652-661.
- Nickel, B. (2010). Ceteris Paribus Laws: Generics and Natural Kinds. *Philosopher's Imprint*, 10(06).
- Pollock, J. L. (1987). Defeasible reasoning. *Cognitive science* 11, 481-518.
- Reiter, R. (1980). A logic for default reasoning. *Artificial intelligence*, 13, 81-132.
- Schroeder, M. (2011). Holism, weight and undercutting. *Noûs*, 45, 328-344.
- Sellars, W. (1953). Inference and Meaning. *Mind*, 62, 313-338.

- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Clarendon Press.
- Thomas, A. (2011). Another particularism: reasons, status and defaults. *Ethical Theory and Moral Practice*, 14, 151-167.
- Travis, C. (2000). *Unshadowed thought*. Cambridge, MA: Harvard University Press.
- Turri, J., W. Buckwalter & P. Blouw. (2015). Knowledge and luck. *Psychonomic bulletin & review*, 22(2), 378-390.
- Väyrynen, P. (2004). Particularism and default reasons. *Ethical Theory and Moral Practice*, 7, 53-79.
- Veltman, F. (1996). Defaults in update semantics. *Journal of philosophical logic*, 25, 221-261.
- Williams, M. (2001). *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford University Press.