

Knowledge Attribution Revisited: A Deflationary Account

Eleonora Cresto

CONICET (National Council for Scientific and Technical Research), Argentina

1. Introduction

For many decades, the concept of epistemic justification has been subjected to endless re-interpretations and re-elaborations, along widely divergent lines (foundationalism, coherentism, internalism, externalism, contextualism, and many others). Regardless of the details, which in many cases are certainly substantive, most mainstream epistemologists seem convinced that the notion of justification has a central role to play in any adequate characterization of knowledge, and hence they tend to think that abandoning justification would be tantamount to resigning the centrality of knowledge in favor of something different – such as true belief. Moreover, according to the usual manner in which we understand how correct knowledge attribution works, we should not attribute knowledge of p to S unless p is true and S is justified in having the belief that p ; this assumption seems to hold even if we shun away from the idea that we can give an

analysis of knowledge in terms of necessary and sufficient conditions.¹

I would like to raise some suspicions on the correctness of this traditional picture. More precisely, I will suggest that justification is not always thought to be a necessary condition for making a true knowledge attribution, according to our pre-theoretical usage of standard epistemic terms. This is not to say that justification is *never* seen as an important requirement; sometimes it certainly is. Still, if the arguments I offer here are on the right track, we can gather that the full-fledged, traditional position on epistemic justification (i.e., the position that justification is a condition on knowledge, and hence on true knowledge attribution) needs to be seriously qualified.² Ultimately, I will contend that this result lends support to a rival epistemological standpoint — what we may dub a moderate Peircean stance on epistemic matters.

Notice that sometimes authors who defend reliability conceptions of knowledge tend to think of themselves as supporting a ‘non-justificationist’ account, to emphasize their rejection of internalist conceptions. But this is of course just a terminological point. My target in this paper is not *internalist* justification, but justification broadly understood – including so-called externalist accounts of justification. To put it differently, all along this paper, by ‘justification’ I will mean a very broad concept, such as ‘whatever it is that an epistemologist would (traditionally) demand from a token of belief to be a piece of knowledge, besides being true’. Thus, ‘being reliably formed’ counts as ‘justified’ in this

¹ For arguments against the possibility of analyzing the concept of knowledge see in particular Williamson (2000); Williamson (2000) does identify several necessary conditions for knowledge, but no *sufficient* conditions.

² Other deflationist accounts on justification can be found in Sartwell (1991, 1992), as well as in Hetherington (2007, 2010). The position I will adopt here, however, is less radical than theirs, and my arguments proceed along very different lines. In addition, people writing within artificial intelligence or epistemic logic tend to equate ‘knowledge’ with ‘true belief’ (cf. Halpern (2005)), although arguably in such cases ‘knowledge’ is in fact a proxy for ‘information’. In any case, part of the goals of the present paper will be to build a bridge between the formal and informal traditions.

broad sense.³

The article is organized as follows. In section 2, I present an example of knowledge attribution to motivate the discussion. I suggest that at the time of evaluating the correctness of the practice of crediting an agent with knowledge that *p* (with the attempt to produce a true statement) we must carefully distinguish between cases where the attributor already believes that *p*, and cases in which she suspends judgment on the truth of *p*. I argue that, in the first case, there are scenarios in which the attribution of true belief is often sufficient to talk about knowledge. I then discuss possible objections to this suggestion. In particular, in section 3 I examine whether justification can be taken to be implicit. This approach can be motivated by the fact that agents are typically disposed to withdraw their knowledge attributions if they come to believe that the agent lacks a suitable justification. In section 4 I propose an explanation for the phenomenon of withdrawal that does not resort to the postulation of implicit justification. As a result of this, we obtain a defense of what I will call a moderate Peircean account. I conclude in section 5.

2. An Example

Norah and Julio meet at a party right after the Senate of their country passes a bill legalizing same-sex marriage, in a marathon session that extended into the night. Norah

³ Someone can contend that, even assuming that sometimes we might be right to attribute knowledge to an agent in the absence of justification, there might still be some extra factor required for knowledge beyond true belief – some extra factor that has nothing to do with justification broadly understood (thanks to an anonymous referee for raising this point). As I am not attempting to provide an analysis of knowledge in terms of necessary and sufficient conditions, this is not a worry I need to address in this paper.

approaches Julio and asks:

“Have you heard the news?”

Norah is certain about the answer already. She is not looking for information about the ongoing events; rather, she wants some information *about Julio*, with the further intention of pursuing a conversation with him about the topic. Julio then says:

“Yes, the Same-Sex Marriage Bill has just been approved!”

Later on that evening Norah is chatting with her friends, the topic comes out again, and one of her friends inquires:

“Does Julio *know* that Egalitarian Marriage is already a law?”

(We can conjecture, for example, that some of Julio’s comments made Norah’s friends suspicious of how well-informed Julio actually was on current politics).

“Yes, he does!” – says Norah, without hesitation.

It sounds like an easy question to her, with a simple, non-problematic answer. At the time of answering it, she is not at all worried about how Julio came to acquire the belief. More generally, she is not at all worried about any other possible property of (what looks to her like) Julio’s particular doxastic structure. It does not occur to her that she might be making a mistake in so attributing knowledge to him, and she is not even aware of the possibility that her answer may be regarded as provisory, pending further inquiry into Julio’s epistemic state; she does not feel that any crucial information is missing in order to be able to answer the question properly.

I take this example to belong to a particular subtype of cases of knowledge attribution; I also want to argue that, once placed in its proper context, it is indicative of a broader phenomenon.

Notice, first, that there are two very different kinds of circumstances in which

we may speculate whether we can correctly assess that *S* knows that *p*:

(I) Either we already believe that *p* is the case, or

(II) We are still in suspense about *p*.

If we are still not convinced that *p* is the case (that is to say, if we are in a type-(II) scenario), then a cautious attitude is typically called for. By judging that *S* knows that *p*, we are simultaneously bound to judge that *p* is true (and not just that *S* believes it to be true); hence, judging that *S* knows amounts to finding a reason to adopt the belief ourselves. We may think of *S* as a reliable source of information, or we may rather want to assess *S*'s evidence by our own lights.

By way of contrast, if we are in a type-(I) scenario, then, as far as we are concerned, *p* is true. Then, how will we judge whether *S* knows? The example examined above seems to tell us that, at least in some cases, *belief attribution alone will suffice*.

Of course, not all type-(I) cases are alike. To begin with, we must concede that knowledge attribution is fallible. Thus, Norah can retract from her attribution if she later finds out that Julio's belief was caused by a particularly vivid dream about the outcome of the debate, after which Julio did not even bother to watch the news. Likewise, if we have prior reasons to doubt Julio's epistemic rationality, we may not be willing to talk about *knowledge* until we gather additional information about the relevant circumstances. Still, the case remains that at least some type-(I) scenarios do not seem to make room for concerns on justification (broadly understood). Let me refer to scenarios that are like the Norah–Julio example in the relevant respects as 'No-Justification Cases' – or, to simplify, 'NJ-cases'.

At this point we may wonder to what extent our practices of knowledge attribution can tell us something important about our pre-theoretic concept of

knowledge. Although some authors expressed reservations,⁴ in general the literature on knowledge attribution agrees upon the fact that we can draw inferences from patterns of usage to concepts – even though, of course, not every possible usage must be taken equally seriously as a means of revealing the nature of a particular concept. In this paper I will take the examples offered to illustrate NJ-cases to be widely representative of a well-established and extended pattern of usage, and, precisely because of that, I will assume that there is *prima facie* a strong presumption in favor of taking usage in NJ-cases as relevant for the concept of knowledge. A word of caution: To say that [there is a strong presumption to the effect that] NJ-cases (which are primarily concerned with knowledge attribution) are relevant for the concept of knowledge does not amount to saying that lack of justification can be automatically read off the examples; after all, there *could* be a good explanation as to why justification seems to be lacking, as we will see below. Rather, what I mean to say is that [there is a strong presumption to the effect that] our best analysis of NJ-cases, whatever that might be, does tell us something about knowledge.

If taken at face value, NJ-cases show that sometimes our pre-theoretic practice is such that we do not care about justification at all, *although we do not resign the use of the concept of knowledge either*. Thus, the scope of our concern for justification issues is more restricted than what has generally been assumed by the epistemological literature. This conclusion may be resisted, however, on several grounds:

(1) It could be objected that even though the example is indeed representative of a standard way of proceeding in a family of cases, in all such cases the standard practice is wrong (perhaps it is even excusable, but still not correct).

⁴ See for example Hazlett (2010), who suggests a “divorce for the linguistic theory of knowledge attributions and traditional epistemology” (p. 500); for a criticism see Turri (2011).

(2) Alternatively, it could be pointed out that, appearances to the contrary, the knowledge attributor in NJ-cases does not fail to demand justification – only, in the absence of conflicting information, the attributor explicitly assumes that the relevant beliefs are justified.

(3) It could be objected that in NJ-cases the justification condition is implicitly assumed to hold: The attributor *would* acknowledge the need to count with a suitable justification, were we to raise the issue explicitly. This would explain, for instance, why the attributor might be inclined to take her assessment back and recognize that the attribution was mistaken, were she to realize that the justification requirement was absent.

Let me try to deactivate these objections.

The first objection goes against a much-entrenched tradition of the literature on the topic, according to which we need to honor our standard usage and intuitions as much as possible. If our theoretical apparatus entails that agents should behave in a way that departs from normal practice, there better be a good explanation for this departure – the epistemologist must provide us with a particular ‘theory of error’.

Perhaps the required explanation can be offered, after all, but I take it that the resulting position will nonetheless be unsettling. NJ-cases are pervasive, and it is likely that our normal understanding of knowledge gets its substance from NJ-cases no less than from more ‘well-behaved’ examples. At any rate, the burden of proof to the effect that this is not the case lies in the justificationist’s quarters.

An obvious problem with the second objection is that this time the analysis does not take the phenomenology of the situation into account: The objection tells us that the attributor does something she does not realize she is doing. Once again, there *may* be a good explanation as to why this is so, but it is far from obvious that this is the right

way to go.

The third objection is probably the most interesting, and will be carefully examined and discussed in the following sections. Before proceeding to the next section, however, let me add a quick clarification. There seem to be two ways of construing objection (3):

(3.a) As claiming that my interpretation of NJ-cases cannot rule out a mainstream explanation of why the attributor would acknowledge the need to count with a suitable justification, were we to raise the issue explicitly; or

(3.b) As claiming that my interpretation of NJ-cases is wrong, because it is *incompatible* with the fact that the attributor would acknowledge the need to count with a suitable justification, were we to raise the issue explicitly.

In this paper I am concerned with (3.b), rather than (3.a). In other words, I take (3) to present an apparent counterexample to my position; I will seek to show that this is not so, and hence that (3.b) is false. I will not presume to prove that no other possible alternative to my explanation of the facts could be true; to put it differently, I acknowledge from the outset that the traditional account of knowledge can *also* provide an explanation of the aforementioned linguistic practice. However, in the end I will argue that my proposal amounts to *a better explanation* than its rivals.

3. Implicit Justification and Withdrawal

As we have seen, the third objection tells us that, appearances to the contrary, in NJ-cases the attributor presupposes the existence of justification; in such examples the attributor's assumption that *S* has justification is simply implicit. Therefore, if the

attributor reflects carefully on what she has just done, she will be forced (i) either to concede that the agent is justified in holding her belief (perhaps by default)⁵; (ii) or to concede that the agent is not so justified, and hence that *she does not know*. Alternatively, the attributor might also conclude (iii) that she cannot determine whether the agent is actually justified or not. In cases (ii) and (iii) the attributor should retract from her attribution: in case (ii) she should admit that her attribution was mistaken, whereas in case (iii) she should admit that her attribution was too quick (because it *could* well be mistaken).

To show that justification is a necessary requirement for true knowledge attribution we need to prove that, in the absence of an adequate justification, the attribution is mistaken. Thus, in what follows I will focus exclusively on cases in which, after reflecting on the topic, the attributor thinks that the agent does not have justification for her belief (as in case (ii) from the previous paragraph) and then withdraws the attribution.⁶

I want to argue that the existence of withdrawals is not sufficient to prove what the objector would like to prove. To put it differently, the fact that agents are typically eager to withdraw prior knowledge attributions under certain circumstances (namely: when they *become aware* of the fact that a justification is lacking) is not equivalent to the idea that in the absence of an adequate justification, the attribution yields a false statement –

⁵ Some authors have proposed that, in the absence of concrete evidence in favor of lack of justification, we can assume that an agent's beliefs are *prima facie* justified; cf. Pryor (2000) or (2004). Alternatively, Michael Williams has argued in favor of an attractive picture of *justification by default*; for this see also Brandom (1994). Thus we read, for example, "...[We] can see justification as exhibiting what Robert Brandom calls a 'default and challenge structure'.... Epistemic entitlement is the default status of a person's beliefs and assertions. One is entitled to a belief or assertion (which, remember, is an implicit knowledge-claim, unless clearly qualified) in the absence of appropriate 'defeaters'; that is, reasons to think that one is *not* so entitled." (Williams (2001), p. 149). In any case, Williams (2001) does not take his view to constitute an objection to the general claim that justification is always a necessary condition for knowledge.

⁶ Case (iii) from the preceding paragraph, in which the attributor suspends judgment on whether the agent has or does not have justification, raises a number of additional problems, but at the end of the day it can be assimilated to case (ii), with only minor modifications.

that is, it is not equivalent to the idea that justification is a necessary condition for true knowledge attribution, i.e., for knowledge. Let us see why.

Objection (3) can be formulated along different lines, depending on whether the objector is or is not committed to rigid standards of justification. If she is so committed – the objection states – the attributor’s willingness to change her prior ascription under certain circumstances can be understood as an explicit recognition that she was mistaken, *precisely* because (as she now realizes) there was no suitable justification for *S*’s belief in the first place. The fact that we tend to describe the previous attribution as a *mistake* tells us that justification is seen as a necessary condition for knowledge.

Alternatively, the objector could seek to explain NJ-cases, as well as possible ulterior retractions, with the aid of a contextualist conception of knowledge,⁷ or perhaps with the aid of other alternatives to contextualism, such as Stanley’s interest-relative invariantism,⁸ or even with some brand of relativism (say, in MacFarlane’s fashion).⁹ Within contextualism, for instance, we might say that NJ-scenarios as described in section 1 involve contexts in which the relevant epistemic standards are just too low. By explicitly focusing her attention on justification issues, the attributor immediately places herself in a different, more demanding context (say, a philosopher’s context), and hence, the standards are bound to go up.¹⁰ Under the new circumstances,

⁷ Such as DeRose (2002) or Cohen (1998). We find a brief discussion of this possibility in Hawthorne (2004), pp. 68 and ss.

⁸ Cf. Stanley (2005).

⁹ Cf. MacFarlane (2005).

¹⁰ It should be recalled here that not all contextualists understand the practice within ‘the philosophy classroom’ along these lines; Michael Williams, for one, would certainly reject this interpretation, insofar as he does not take the skeptic to raise the standards, but to change the subject (2001, Chapter 16).

Regardless of the details on justification standards, the structure of objection (3) goes as follows. Let us concede, for the sake of the argument, that this premise is true:

Objection (3) boils down to the claim that, given (WITHDRAWAL), the absence of justification for p (by S) makes it false that S knows that p :

or, to use a shortcut:

where ' J_{sp} ' stands for 'S is justified in holding that p ' (according to *some* standard) and ' K_{sp} ' stands for 'S knows that p '.

11

(b) If S^* is/becomes aware of the fact that S does not have justification for p , then S^* is (typically) not ready to attribute knowledge of p to S .

The consequent is meant to be read descriptively (it tells us what S^* is typically disposed/not disposed to do); this is to say that (b) as presented here is an empirical claim on usage, though one may eventually aim to use it as a means to discover conceptual facts. But even granting the correctness of the path from usage to concepts in this particular case (as discussed on section 2), all we can conclude from (WITHDRAWAL) is:

(c) If S^* is/becomes aware of the fact that S does not have justification for p , then S^* cannot truthfully attribute knowledge of p to S ;

where, of course, to say that S^* cannot truthfully attribute knowledge of p to S amounts to saying that in such circumstances S does not know that p :

(d) If S^* is/becomes aware of the fact that S does not have justification for p , then S does not know that p ;

more concisely,

(d') $B_{S^*} \sim JSp \text{ } \mathcal{A} \text{ } \sim K_S p$

Subscript ' S^* ' refers to the attributor, and ' B ' stands for the operator of being or becoming aware of a certain state of affairs; in many contexts ' B ' could just be replaced by an operator of true belief, insofar as we understand it either as 'occurrent belief', or as a reflective state of some kind – as opposed to, say, a mere disposition to act in a certain way.

Let me put it in a slightly different manner. We can agree that (WITHDRAWAL) is indeed true: let us call it our *evidence*. Our evidence tells us that being aware of the lack

of justification (or perhaps even being consciously in doubt as to whether justification is indeed lacking) prevents us from regarding ‘*S* knows that *p*’ as a true statement. Now, traditional epistemology proposes an *explanation* for this evidence. The explanation is that we take justification to be necessary for ‘*S* knows that *p*’ to be true. However, as it is often the case, this explanation is not logically entailed by the evidence, and, as it happens, we can indeed come up with a *rival* one – at which point we can start to discuss which of the two is actually better.

According to (d), then, what is necessary for true knowledge attribution is *not to consciously (and truthfully) believe that *p* is not justified for *S** (according to our favorite standard). This is not equivalent to saying that justification is a necessary condition. The crucial point here is that, unlike (a), (d) licenses a framework according to which, at times, certain elements of *S*’s antecedent doxastic state *are not the type of entity which can have or lack a justification, in a strong, full-fledged sense*. This can happen, in particular, when the elements under consideration are *also* part of the attributor’s doxastic background (such as, in our example, “The Senate approved the Equal Marriage Bill”). If this diagnosis is correct, simply assuming that such elements are always either justified or unjustified in a strong sense would be akin to committing a sort of category mistake.

In order to capture this suggestion, it may prove fruitful to distinguish between *not being justified* (or *not having justification*),¹¹ on the one hand, and *being unjustified*, on the other. As before, let ‘ $\sim JSp$ ’ stand for ‘*p* is not justified for *S*’; in addition, let ‘*p* is unjustified for *S*’ be rendered as ‘ U_{Sp} ’. $\sim J_{Sp}$ is meant to be compatible with both U_{Sp}

¹¹ One can argue that there is a difference between stating that *p* is justified for *S* (at *t*), and stating that *S* has justification for *p* (at *t*). However, the difference (if indeed there is one) is not relevant for the present context of discussion, so I will gloss over it.

and $\sim U_S p$, while $\sim U_S p$ is meant to be compatible with $J_S p$ and $\sim J_S p$; on the other hand, $J_S p$ entails $\sim U_S p$.¹² I will assume, moreover, that if an agent becomes aware [believes truthfully and consciously] that p is not justified for S , then p is unjustified for S :

$$(e) B_{S*} \sim J_S \text{ } \mathcal{A}E \text{ } U_S p$$

Now, we shall see that even if we take the conditional

$$(f) U_S p \text{ } \mathcal{A}E \text{ } \sim K_S p$$

to be correct, the conditional

$$(g) \sim J_S p \text{ } \mathcal{A}E \text{ } \sim K_S p$$

need not be: as it will become apparent soon, p could lack justification just in the sense that p is not in the appropriate context for justification to arise in the first place. I will articulate this proposal with more detail in the next section.

4. A Moderate Peircean Proposal

In “The Fixation of Belief”¹³ Peirce famously contrasted real, living doubts to ‘paper doubts,’ or ‘philosophical doubts’. As opposed to the latter, real doubts have a clear, distinct effect on behavior, which also differs from the effect we expect from beliefs. Real doubts make us uncomfortable, and thus compel us to engage in research; the aim of research is primarily to settle our mind, rather than to obtain the truth – even though, as it happens, we only settle for what we think is true (and, indeed, we think that each of our

¹² See footnote 18 for a more detailed account of how these relations go.

¹³ Peirce (1877).

beliefs is true). When doubt ceases, so does the impulse that makes research possible (and meaningful) in the first place. Thus Peirce offers a radically anti-Cartesian epistemology. Notice, in particular, that Peirce's position entails that we should not indulge in skeptical worries: philosophical skepticism is simply not to be paid attention to. Even more generally, his position entails that agents should not seek to have, as a starting point, a set of indubitable propositions, but only a set of propositions that are free from *real* doubt – that is, propositions about which agents are not, as a matter of fact, in doubt. In his mature writings, these seminal ideas led him to build an interesting and complex framework of how inquiry proceeds (his famous sequence of abduction-deduction-induction), which I cannot dwell upon here.

Regardless of the nuances of Peirce's writings, or of his exact intentions, ever since the 1970s some writers have conceived of themselves as following a Peircean model, in a broad sense. It has been baptized 'the Belief-Doubt Model', by opposition to the so-called 'Doubt-Belief Model', of Cartesians roots.¹⁴ According to this perspective, when we do research we never start from scratch, so to speak, but we always find ourselves holding a particular belief corpus which is, as far as we are concerned, true. In the course of our investigation sometimes we feel compelled to revise our corpus (say, as the result of new observational inputs, or as a result of being exposed to conflicting ideas). In this case – but only in this case – justification is called for: justification applies to *changes*, exclusively.¹⁵ The Belief-Doubt Model emerges as an alternative to both foundationalism and coherentism. This neat way of framing the task of epistemology has led to the

¹⁴ Cf. Levi (1997), Ch.1, and Fuhrman (1997), pp. 4 and ss., among other places.

¹⁵ On this also see Bilgrami (2000).

development of sophisticated theories of belief dynamics, which in some cases bear much in common with Bayesian-based programs.¹⁶

Here I do not want to commit myself to a full-fledged Peircean framework. I am ready to concede that it makes sense, at times, to ask whether particular antecedent beliefs are justified (as opposed to insisting that justification applies to epistemic changes only). However, talking about antecedent justification requires a good deal of caution. A suitable account of true knowledge attribution should tell us precisely under which circumstances justification becomes important. Let me elaborate on this point a bit. I will dub the resulting position ‘a moderate Peircean standpoint’.

We should start by noticing that the concept of knowledge applies normally to two different kinds of situation. On the one hand, as agents naturally engaged in inquiry, we are interested in evaluating the reliability of sources that may lead to profitable changes of mind. In particular, the act of assessing the linguistic emissions of a fellow agent as *knowledge* has a potential future utility to us. But, on the other hand, we are bound to judge the world from our very point of view. In this sense we are the measuring rod, so to speak: we trust our own judgment, in whatever it is that we trust it (in whatever it is that we are subjectively certain of); this is of course tautological, but still worth reflecting upon. Putting the two ideas together, we use ‘knowledge’ both to refer to something we assume to have, and to something we hope to obtain.

In the light of this, the concept of justification is useful insofar as it helps us obtain new true beliefs; those beliefs we already have and about which we feel (for the time being) certain are not subjected to justification qualms (because we already have them!) —

¹⁶ See paradigmatically Levi (1980). Part of the belief revision tradition in philosophy and in AI has been conceptually inspired by the Belief–Doubt Peircean model. For a recent survey on the Belief Revision Program see Pederson and Arló-Costa (2011).

except when we come to consider the possibility of revising. In other words, the concept of justification only applies to antecedent beliefs under certain conditions, namely, *when we think of ourselves as potentially involved in a revision process* concerning the beliefs under consideration. Moreover, according to this picture our use of the concept to characterize beliefs of a third party agent is derivative from its usefulness as captured by the first person point of view; which is to say that the semantics of ‘justification’ and ‘belief’ follow this structure:

- (h) An antecedent belief [by *S*] acquires the ability to be justified only when seen as a potential candidate for acceptance/rejection [by some agent *S**, who may or may not coincide with *S*].

Or, more generally:

- (i) *x* acquires the ability to have *y* only under *z*.

Notice that this general structure is not peculiar to belief and justification; we can find analogies everywhere. Consider:

- (j) Kids acquire the ability to be good students only when they start school.
- (k) Viruses acquire the ability to replicate themselves in an efficient manner only when they enter a particular living cell.

It does not make sense to assert that the child either was or was not a good student beforehand; likewise, the efficiency of virus replication depends crucially on the medium: it does not make sense to say that they are either efficient or inefficient at replication beforehand (and independently of a specific living environment).

On a closer look, it can be argued that the finer structure behind all such cases actually goes like this:

- (l) *x* cannot be a good *z* before becoming a *z*,

which gives us:

- (m) A belief cannot be a *good* candidate for acceptance/rejection before becoming a candidate for acceptance/rejection in the first place.¹⁷

As it happens, a belief becomes such a candidate only upon reflection. Thus, we can go one step ahead and claim that *reflection is the condition of possibility of justification acquisition*; the possibility of being justified does not pre-exist before a reflective state takes place.¹⁸

In any case, we should better examine carefully what ‘reflection’ exactly amounts to in this discussion. Reflection *per se* is not sufficient for the conceptual space of justification to arise. We can reflect all we want on whether someone told us that *p*, and yet not reflect on whether that person was *entitled* to tell us that *p*. The first type of reflection does not deal with justification issues at all. In the original Norah–Julio example, in order to judge that Julio knew that Egalitarian Marriage was already a law, Norah was bound to be in a reflective mood, at least in some sense, even if she did not take justification issues into account. So the relevant concept here is *reflection with a specific goal in mind* — namely, the goal of a potential belief change. Let me call it ‘J-

¹⁷ Where a belief is a ‘good’ candidate for acceptance/rejection if it is a belief a rational epistemic agent could accept/reject (i.e., if it is a belief one would be justified in accepting/rejecting).

¹⁸ In the light of this, here is a more detailed account of the formal tool used to capture the distinction between *not having justification* and *being unjustified*. A central assumption of Moderate Peircean Epistemology is that *bona fide* justification entails that some agent has reflected upon the proposition under consideration. Hence, we start by assuming:

(1) $J_S p \text{ } \mathcal{A}E \text{ } R_{S^*} p$

where ‘ $R_{S^*} p$ ’ stands for ‘ S^* reflected on p ’. From (1) it follows:

(2) $J_S p \text{ } \mathcal{A} \text{ } (J_S p \text{ } \& \text{ } R_{S^*} p)$, and therefore

(3) $\sim J_S p \text{ } \mathcal{A} \text{ } ((R_{S^*} p \text{ } \& \text{ } \sim J_S p) \vee \sim R_{S^*} p)$

Now we can define:

(4) $U_S p = \text{df. } (R_{S^*} p \text{ } \& \text{ } \sim J_S p)$, from which we obtain

(5) $J_S p \text{ } \mathcal{A}E \text{ } \sim U_S p$

as well as all the compatibility relations between *J* and *U* discussed in section 3.

Reflection'. As opposed to a standard reflective stance, J- reflection can be said to assume a particular *deliberative* perspective on our doxastic state, rather than a merely contemplative, or judgmental, perspective.

A more detailed elaboration of this proposal is beyond the limits of this paper; for the moment, let me mention a few consequences. First, there is room to contend that, if true knowledge attribution demands justification only when we adopt a J-reflective stance, then there is a sense in which the epistemic contextualist is right, and a sense in which he is not. He might be right in claiming that the meaning of 'knowledge' changes when we enter the philosophy classroom; he is not right in claiming that the switch reflects a change in justification standards.¹⁹ Rather, the change amounts to the creation of a conceptual space in which justification talk becomes meaningful.

Second, the distinction between a reflective stance, broadly construed, and a more specific realm of J-reflection is also relevant for possible discussions concerning the link between knowledge and epistemic *responsibility*, particularly within the context of self-attribution of knowledge. For example, consider a theory according to which ideally responsible agents are capable of reflecting on their own mental states and making them their own, so to speak, by means of second-order attitudes (the slogan here could be: 'make sure you own your own beliefs'). At this point someone might wonder whether responsibility so understood is not but a side effect of an internalist conception of justification. It is not: responsible agents need not satisfy justification in the internalist sense. Reflection in the sense the internalist demands us to do is sufficient, but not *necessary*, for being a responsible agent. We can attain responsibility by reflecting on our beliefs without J-reflecting on them; if they are true, they might still amount to knowledge.

¹⁹ But see footnote 10.

Thus, by emphasizing that not every reflection stance is simultaneously a justification stance, we free the concept of epistemic responsibility, as captured by second-order attitudes, from its historical submission to an internalist epistemology.

Third, notice that, within the present account, a belief can only be justified if someone reflects on whether it is indeed justified. A possible worry here is whether we are turning justification into an extremely subjectivist concept. However, this worry is unfounded. It is true that, at a given time t , a proposition p can belong to X 's conceptual space for justification, but not to Y 's – hence we can say that the truth-value of assertions expressing justification are partly relative to the attributor. But they are relative only in the following sense: what is up to the attributor is the truth-value of the disjunction (either p is justified or strongly unjustified for an agent S). Once p is in S^* 's conceptual space for justification, it is *not* up to S^* whether S is or is not justified in having p . For example, an externalist on justification can still maintain that, *provided* p is already in the realm of justification, then p is justified only if certain objective relations with the world obtain, regardless of the agent's awareness of the existence of such relations. In other words, I am assuming that there is a fact of the matter as to whether S^* is or is not right in claiming that S is or is not justified (regardless of which fact of the matter that may be). Thus, S^* may well be wrong as to whether p is or is not justified for S (where S may or may not be S^* itself). In short, knowledge attribution does not require our being in a J-reflective mood; however, once we locate ourselves in the realm of justification, justification can consist of whatever it is that our favorite theory tells us about it.

In light of this, let me stress, then, that by construction of the example, at the time of withdrawing her attribution what Norah discovers is not just that Julio's belief lacks

justification: she discovers that the belief was *unjustified*. Recall that it is Julia's reflection on the topic what creates the possibility space that enables justification talk in this case. Once such a space has been created, it would be conceptually impossible for Norah to 'discover' that Julio is, say, not justified but at the same time not unjustified. Notice, moreover, that discovering that Julio himself had not reflected on the justification problem is not sufficient to say that Julio is unjustified; actually, if Norah is an externalist, Julio need not so engage in reflection for his belief to be justified in the first place.

Let me further elaborate on this point. To what extent is Moderate Peirceanism truly compatible with externalism? At this stage a reliabilist could object that, if at all, Julio's belief about the law has already been reliably formed when Norah took it to be knowledge, before raising any questions about its justification – still Moderate Peirceanism would claim that in that case the belief was not justified.²⁰ The answer is that for Moderate Peirceanism, being reliabilists about justification means that we take a belief to be justified if, *assuming it is already in the space of justification*, it has been reliably formed. So it is correct to say that a reliably formed belief is not yet a justified belief until someone reflects on whether it is reliably formed. But this does not commit us to internalism, at least not in the standard sense of this word. To begin with, justification requires reflection *by someone*, but not necessarily by the very subject who holds the belief. In addition, the person who so reflects may well be unable to tell whether the belief is indeed reliably formed or not; as long she raises the problem of justification, if the belief has been reliably formed, then it is justified, regardless of whether someone (let alone its owner) is aware of its having been reliably formed. In short, an epistemologist committed to Moderate Peirceanism who also

²⁰ Thanks to an anonymous referee for raising this objection.

shares reliabilist intuitions cannot be labeled as ‘internalist’ in any interesting sense. I concede that if one gets too strict on the ‘externalist’ label one might want to refuse to apply it to such a position either. But in this case, what we should say is that Moderate Peirceanism turns the distinction between internalism and externalism into a false dichotomy. I would be happy to accept this consequence.²¹

Fourth, the fact that justification requires a J-reflective stance has direct consequences on our understanding of knowledge. According to Moderate Peirceanism, knowledge of p demands that p not be unjustified for S , but not being unjustified is relative to a certain epistemic perspective. The upshot is that “ S knows that p ” is in fact an abbreviation for “ S knows that p within S^* ’s overall epistemic perspective”. In other words, the truth of knowledge attributions are relative to certain epistemic perspective. Notice, however, that this is not tantamount to saying that the truth-value of “ S knows that p ” is up to S^* , as S^* could get things wrong (just as the truth-value of “ p is justified for S ” is not up to S^* , as we have seen). For example, if p is in fact false, but S^* takes it to be true, S^* might mistakenly think that S knows. Alternatively, if S^* is an externalist and she is mistaken in thinking that p has been reliably formed, then again S^* might wrongly believe that S knows. The reason is that, although S^* ’s epistemic perspective can make a difference on whether S ’s beliefs are or are not in a justification space, by contrast S^* ’s epistemic perspective typically does not affect whether p is or is not justified provided p is already in a

²¹ Moreover, consider what happens if S^* is wrong in (consciously) believing that p is not justified for S . In this case it makes sense to say that, even though S^* is not ready to attribute knowledge to S , perhaps S^* *should* be so ready. In other words, if S^* falsely believes that p is not justified for S , it is not true that p is unjustified for S in the strong sense. Fortunately, ‘being or becoming aware of φ ’ (as captured by operator ‘ B ’ from the previous section) is factive, so we do not need to worry about this problem. In any case, regardless of how wrong S^* may be, to the extent that S^* has considered the topic, p can no longer be neither justified nor strongly unjustified for S , as I have already pointed out.

justification space, nor does it affect the truth-value of p .²²

Note that a moderate Peircean standpoint is sufficient to account for the problems discussed in the previous sections. First of all, it is easy to see that it provides a straightforward explanation of NJ-cases *by taking the linguistic evidence at face value*. This is certainly more than what its rivals can do. Moreover, notice that (WITHDRAWAL) can also get a natural explanation within the present epistemological framework. When we engage in J-reflection we are in a peculiar state of mind: the state of mind in which we *put ourselves in the agent's shoes*, as it were, *and foresee a potential revision* — and that is precisely why justification becomes important. It does not follow, however, that an interest in justification is implicit beforehand; a deliberation mood is not the *default* mood.

Borrowing from the terminology of the previous section, the linguistic evidence is compatible with the following picture: being unjustified prevents S from having knowledge, while merely not being justified does not. In other words, the conditional

$$(n) U_{Sp} \text{ } \mathcal{A} \text{ } \sim K_{Sp}$$

is correct, whereas

$$(o) \sim J_{Sp} \text{ } \mathcal{A} \text{ } \sim K_{Sp}$$

²² Thus, for example, suppose that Juan formed a particular belief as a result of taking a pill, and suppose the belief happens to be true, as a matter of luck. Could it count as knowledge? (Thanks to David Etlin for raising this question). If I do not know anything about the fact that Juan took that pill, and Juan tells me something true in which I already believe (and suppose I do not have further reasons to doubt Juan's epistemic credentials, generally speaking) then it is quite natural for me to say that *Juan knows what he just told me*. According to Moderate Peirceanism, in such circumstances I would have uttered a true sentence. At this point the objector would jump in and claim, "Oh, but Juan doesn't *really* know it, does he?" All this shows is that it is of course correct to say that Juan does not know, *from the vantage point of view of those who know that Juan took the pill*. Once we are in a J-reflective stance, we cannot avoid passing judgment on S 's justification status regarding p , and hence we cannot avoid thinking that the knowledge attribution was a mistake, if p happened to be unjustified.

is not.

Someone could object that such an explanation does not come naturally; on the contrary – so the objection goes – we are naturally inclined to understand (WITHDRAWAL) as a symptom that justification was implicit. As I see it, however, it is far from clear whether this attitude is indeed ‘natural’. It might well be that it is rather our philosophical education which has trained us to think along these lines. In other words, we can very well suspect that the philosophical interpretation of (WITHDRAWAL) is often biased in favor of traditional accounts of justification, and this is what generates the appearance of ‘naturalness’ for the more standard explanation. A final verdict on this matter should be left to experimental philosophy. For the moment, I rest content with having put forward a reasonable doubt on the claim that postulating implicit justifications is always the more obvious option for normal speakers.

In light of the previous discussion, we have grounds to claim that Moderate Peirceanism amounts to a better explanation of the facts. It is true that both Moderate Peirceanism and traditional accounts are compatible with the data; as it is often the case, here we find underdetermination of theory by data. In this paper I started by revising certain linguistic evidence (the NJ-cases). As it happens, an interpretation of the evidence always requires a particular framework. I cannot give a knock-down argument as to why we should choose the new framework, but I can give a reason to think that the alternative explanation is actually *better*: to wit, I can show that the alternative explanation is less costly, in several aspects.

Why is it less costly? For the traditional framework, the initial knowledge attribution (in the context of NJ-cases) is interpreted as pointing to the fact that justification

is implicit; in order to get support for this claim we need yet another piece of evidence: Withdrawal. And Withdrawal, in turn, is interpreted as pointing to the fact that justification is a necessary condition for knowledge.

For the Moderate Peircean framework, by contrast, the initial knowledge attribution gets explained in a straightforward way: no additional (implicit) justification claim is needed. Then there is another piece of evidence we need to accommodate: Withdrawal. Withdrawal, in turn, is easily explained by claiming that, in order to make a correct knowledge attribution, justification cannot be explicitly absent once the worry about justification issues is raised. Notice that the traditional framework is *also* committed to this last claim. Only, *in addition to it*, the traditional framework also claims that, in order to make a correct knowledge attribution, justification cannot be absent *even when we are not playing the justification game*. From the point of view of Moderate Peirceanism, this last bit does not make much sense. In short, the traditional framework is forced to make additional claims to explain Withdrawal, and, in addition, it cannot take the initial evidence from NJ-cases at face value.

There is a final point worth discussing. Moderate Peirceanism is not meant to contest the correctness of the truth or belief conditions endorsed by the traditional picture. As with justification, agents are typically ready to withdraw knowledge attributions if they realize that belief or truth is not present. However, in such cases I am ready to say that the traditional account is correct in asserting that both truth and belief claims are indeed *implicit*, as I do not have any other better explanation for this fact. For example, whenever we listen to the words of any agent *S*, we tend to attribute beliefs to *S* that match those words (by contrast, we would need some extra cues to think *S* was not being sincere); thus, in our example, Norah attributes a belief to Julio in a completely unproblematic way, and

so do we.

5. Conclusions

I have outlined an account according to which justification is not a general requirement for true knowledge attribution. Moderate Peircean Epistemology tells us that we should qualify the traditional standpoint on justification for knowledge attribution, while acknowledging the importance of justification for particular circumstances. Let me stress here that Moderate Peirceanism does *not* say that the concept of justification is not relevant in itself, or that it should be abandoned. *If* the question for justification arises – which is actually like saying: “if we adopt a deliberative perspective towards a particular set of beliefs, or if we conceive of a third party agent as potentially engaged in such a deliberative perspective” – then we are reluctant to talk about knowledge unless the belief is found to be justified. However, sometimes justification questions just do not come up. No matter how often we engage in deliberation, it is not the natural attitude towards our beliefs – not even towards those beliefs we do not hesitate to call ‘knowledge’, pre-theoretically speaking (say, among competent speakers of the language who need not have taken epistemology courses).

As we have seen, the proposed standpoint gets straightforward support from NJ- cases – although, of course, NJ-cases by themselves are not sufficient to dislodge once and for all more traditional perspectives. Why should we favor the present proposal over the traditional picture, then? Let me summarize some of the considerations discussed in previous sections.

To begin with, if we pay due attention to the linguistic evidence we notice that the normal discourse of a competent speaker is sometimes neutral, and sometimes plainly hostile to the idea that we need to assume the satisfaction of the justification requirement to be able to make correct attributions of knowledge. NJ-cases constitute a prime illustration of this situation, and apparent counterexamples such as (WITHDRAWAL) can be explained within the same perspective. Thus, an advantage of the present proposal is that it can take the linguistic evidence from NJ-cases at face value. As a result of this we obtain a simpler account; we are not forced to postulate the existence of implicit justification to restore consistency between theory and data.

In addition, the present proposal points to an attractive picture of the different goals we may pursue at the time of attributing knowledge. By distinguishing between a deliberative and a judgmental type of reflection, we can put the concern for epistemic revisions at the center stage, while also making room for the possibility of antecedent justified beliefs, under certain restricted conditions. In this way we can reconcile the gist of Peircean-based intuitions with more traditional intuitions of mainstream epistemology; thus we can get the best of the two worlds, so to speak.

Let me end by commenting on the relevance of obtaining a partial vindication of a Peircean framework of sorts. In general, contemporary elaborations of the Peircean picture do not devote much effort to actually argue in favor of the advantages of embracing a Belief-Doubt model. Typically, the reader is just invited to explore by herself how far she can go if she switches to a younger, still not-too-well-explored

epistemological standpoint.²³ The present work can be seen as an attempt to raise interest in some Peircean themes, by connecting them with the worries and discussions of a wider audience.

Acknowledgments

Previous versions of this paper have been presented at *The First Congress of the Latin American Association for Analytic Philosophy* (Mérida, Yucatán, April 2010); at *The XV Congress of the Argentine Philosophical Association* (Buenos Aires, December 2010), and at *SADAF International Workshop 2011* (Buenos Aires, August 2011). I am thankful to the participant for their questions and suggestions.

References

- Pedersen, A. P. and Arló-Costa, H. (2011). "Belief Revision". In L. Horsten and R. Pettigrew (eds.), *Continuum Companion to Philosophical Logic*, London and New York, Continuum Press, pp. 450-502.
- Bilgrami, A. (2000). "Is Truth a Goal of Inquiry? Rorty and Davidson on Truth." In R. Brandom (Ed.), *Rorty and his Critics*. Oxford: Blackwell, pp. 242-261.
- Brandom, R. (1994). *Making it Explicit*. Cambridge, Mass., Cambridge University Press.
- Cohen, S. (1988). "How to be a Fallibilist." *Philosophical Perspectives*, 2: 91-123.

²³ In Isaac Levi's terms, the traditional epistemological framework would amount to a degenerate research program, in Lakatosian terms: almost 400 years later, we are still running in circles, wondering how to overcome evil demons of various types. See for instance Levi (1997).

- DeRose, K. (2002). "Assertion, Knowledge, and Context." *Philosophical Review*, 111: 167-203.
- Fuhrman, A. (1997). *An Essay on Contraction*. Stanford: CSLI Publications.
- Halpern, J. (2005). *Reasoning about Uncertainty*. Cambridge: The MIT Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hazlett, A. (2010). "The Myth of Factive Verbs." *Philosophy and phenomenological research*, 80: 497–522.
- Hetherington, S. (2007). "Is This a World Where Knowledge Has to Include Justification?" *Philosophy and Phenomenological Research*, 75: 41-69.
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge: On Two Dogmas in Epistemology*. Oxford University Press, Oxford.
- Levi, I. (1980). *The Enterprise of Knowledge*. Cambridge, MA: The MIT Press.
- Levi, I. (1997). *The Covenant of Reason*. Cambridge: Cambridge University Press.
- MacFarlane, J. (2005). "The Assessment Sensitivity of Knowledge-Attributions". In T. Gendler-Szabo and J. Hawthorne (eds), *The Oxford Guide to Epistemology*.
- Peirce, C.S. (1877). "The Fixation of Belief". In Charles S. Peirce (1992). *The Essential Peirce*, Vol. I. Edited by Nathan Houser and Christian Kloesel. Bloomington and Indianapolis: Indiana University Press.
- Pryor, J. (2000). "The Skeptic and the Dogmatist." *Nous*, 34: 517-549.
- Pryor, J. (2004): "What's Wrong with Moore's Argument?" *Philosophical Issues*, 14: 349-378.
- Sartwell, C. (1991). "Knowledge is Merely True Belief". *American Philosophical Quarterly*, 28: 157-66.
- Sartwell, C. (1992). "Why Knowledge is Merely True Belief." *The Journal of*

Philosophy 89: 167-80.

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Turri, J. (2011). "Mythology of the Factive". *Logos and Episteme*, 2: 143:152.

Williams, M. (2001). *Problems of Knowledge. A Critical Introduction*. Oxford: Oxford University Press.

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.