

On Reasons and Epistemic Rationality*

Eleonora Cresto

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

[Penultimate version. Last version to appear in *The Journal of Philosophy*]

In a recent paper, John Hawthorne and Jason Stanley propose an analysis of how knowledge and action relate to each other.¹ According to their *Reason-Knowledge Principle* (RKP), it is appropriate to treat the proposition that p as a reason for acting iff we know that p , for p -dependent choices;² within their account, in addition, knowledge delivers probability 1. Hawthorne and Stanley also note that sometimes it is intuitively rational to act on partial beliefs. What is appropriate to treat as one's reason for action, in this case, is the epistemic probability of p conditional on the agent's total knowledge K . RKP then requires that one knows that $P(p | K) = r$ (for some r).³

Here I am not concerned with the general formulation of RKP, but I shall focus exclusively on the more restricted contention according to which agents should not invoke probability claims as reasons unless they know that such claims are true. I believe there are grounds to think that this is a problematic demand.

* A previous version of this commentary was read at a workshop held by the *Grupo de Acción Filosófica* (GAF) at the *Universidad de Buenos Aires*, in April 2008. I am indebted to Jason Stanley for valuable feedback. I also want to thank the members of GAF for discussion and comments.

¹ Hawthorne and Stanley, "Knowledge and Action," this JOURNAL, CV, 10 (October 2008): 571-90.

² A choice between options $x_1...x_n$ is said to be p -dependent iff the most preferable of $x_1...x_n$ conditional on the proposition that p is not the same as the most preferable of $x_1...x_n$ conditional on the proposition that not- p .

³ This is my notation, not theirs.

I would like to begin by reflecting on the role of personal probabilities at the time of justifying action. The authors make it clear that RKP is meant to refer to *objective* probability functions. Still, they do not say explicitly whether, in their view, subjective measures can never act as motivating reasons on their own right. At any rate, the linguistic evidence does not seem to exclude this possibility in any obvious way, as I hope to show below. In particular, note that the fact that a given probability claim is best interpreted as being epistemic does not mean it is not subjective, where subjective measures may very well incorporate estimates about (physical) chances. Unfortunately, if probability judgments are – at least at times – taken to encode personal measures, then RKP falls short of what we need.

Indeed, Hawthorne and Stanley concede that RKP does not mesh well with personal probabilities (584). But the actual explanation as to why this is so is not addressed by their paper. The crucial point is that positing knowledge – or even belief – adds an unnecessary complication, and ultimately distorts the nature of the underlying phenomenon. In this respect, there seems to be an interesting analogy between credences and desires, preferences, possibility judgments, or aesthetic judgments. Suppose that, as far as *S* is concerned,

x is desirable;

it is correct to do *y*;

it would be nice if *p* obtained;

p is preferable to *q*;

p is possible;

p is highly probable.

In each case we can identify a primary attitude that consists in desiring a particular object, or preferring the occurrence of a particular state of affairs, or conceiving of the occurrence of a particular state of affairs as more or less probable. All such attitudes play a vital role in the economy of an agent's epistemic life, at the time of engaging in both theoretical and practical reasoning. Suppose now that *S* believes (knows) that she is committed to a particular set of personal probability judgments (desires, judgments of taste, etc.). At least in typical scenarios, the corresponding second-order belief (knowledge) will not do any real work in *S*'s acting or reasoning in a certain way, over and above what is already achieved by the first-order level.⁴ Consider, by way of illustration: “Why have you moved your arm?” “Because I wanted to reach the bottle and drink some water” (rather than: because I believed/ knew that I wanted to reach the bottle and drink some water); “Why have you bought that paint?” “Because I like it” (rather than: because I believe/ know that I like it); “Why are you carrying an umbrella with you?” “Because it seems likely [to me] that it will rain” (rather than: because I believe/ know that it seems likely that it will rain). Thus, if we are to trust our ordinary use of the language, the reason I moved my arm was a primary desire, so to speak, and not a second-order belief, or a piece of second-order knowledge, about my having a particular desire.⁵ Likewise, the reason I am carrying an umbrella is a primary probabilistic commitment, and not a belief, or a piece of knowledge, about a particular probabilistic claim.

Notice, moreover, that first-order knowledge claims behave very much unlike putative second-order propositional attitudes on personal probability judgments. “I

⁴ By the expression ‘second-order belief’ I mean to refer to a belief about a first-order attitude that may, but need not, be itself a belief; *mutatis mutandis* for ‘second-order knowledge’.

⁵ Of course, this is not to deny that a complete account of my reasons for acting may well incorporate, in addition, an array of *first-order* (full) beliefs, as well as further desires. (Thanks to Alejandro Cassini for pressing this point).

invested in the market because the profit was going to be high” and “I invested in the market because I knew the profit was going to be high” are often interchangeable; “I invested in the market because it seemed very likely to me that the profit was going to be high” and “I invested in the market because I knew it seemed very likely to me that the profit was going to be high” are not – the last assertion is just awkward. Its awkwardness tells us something important about how partial beliefs enter into the business of giving and asking for reasons.⁶ Once we admit personal probabilities into the picture (and it is not clear what would prevent us from doing so) motivating reasons might turn out to be other than the content of (true, justified) beliefs: primary probabilistic commitments can also do the trick. Thus, it seems perfectly right to treat a particular probabilistic commitment *C* as a reason for acting, without thereby requiring knowledge of *C* – or belief, for that matter.

It might be objected here that having *C* requires our knowing that we have it, out of rationality considerations. But this point is irrelevant for the present discussion: regardless of the intrinsic value of second-order attitudes, the examples presented above show that our reasons for acting can typically be found in primary commitments – epistemic and otherwise. To put it differently: if an agent has *C* but does not know she has it (say, because of transparency failure), she is already at fault; her further treating *C* as a reason does not add any extra offence. In short, if probabilistic talk is interpreted along subjectivist lines, RKP can be violated without intuitively making the agent accountable as far as her treatment of reasons is concerned.⁷

⁶ Of course, we could always conceive of particular scenarios in which focusing on second-order attitudes becomes acceptable (“Are you sure you don’t want to try the cake?” “Yes – I *know* I don’t like chocolate”). But this is beside the point – the fact remains that mentioning second-order attitudes on probabilistic commitments, desires or preferences is usually idle, and it very often leads to infelicities.

⁷ Incidentally, to say that RKP *can* be violated is not equivalent to saying that our failure to know that we have *C* should never prevent our treating *C* as a legitimate reason; we might still contend, for instance, that treating *C* as a reason should imply that it is *true* that we have *C* (say, if we understand ‘reasons’ in

Let me turn now to Hawthorne and Stanley's preferred interpretation of probabilities. A theory of objective functions of the sort required by the authors would need to tell us how to obtain objective confirmation measures between p and K , for any possible p and K . But we are not given any hints as to how such a confirmation theory could go; more importantly, we are not given any reassurance that such a theory is possible in the first place. In the absence of any details, we seem to be left with how much S takes K to confirm p – but this, of course, takes us back to the realm of personal measures.

In any case, the authors' idea is that, ultimately, by focusing on the evidence we can circumvent mentioning probabilities altogether: acting on knowledge of epistemic probabilities would be tantamount to acting on those propositions conditional on which we define the particular epistemic measures we have (584-585). To know that $P(p | K) = r$ is just to have K . If this were true, we would indeed get rid of the problem of providing a suitable interpretation for probability-talk in the natural language, at least vis-à-vis an analysis of the link between knowledge and reasons.

But this move will not do. Consider an agent who asserts,

- (*) “The reason I treated patient A with drug d on this occasion was that, as far as I know, drug d cured some people in the past, killed others, and, moreover, several other untreated patients with symptoms similar to those of A died a horrible death.”

the same objective way Hawthorne and Stanley do). In any event, notice that the situation here is not analogous to the one the authors have in mind when they present their principle in terms of full beliefs. Unlike the case in which an agent falsely believes a given proposition p about the external world, the agent who misidentifies her probabilistic commitments can be charged with irrationality, rather than with a mere factual mistake. Hence, once again, there is room to argue that treating C as a reason does not add further irrationality.

Is this an admissible way for the agent to justify her action? Hardly so; we just cannot see where the motivation for her behavior lies. The awkwardness of the agent's discourse reveals precisely that we cannot assume probabilities to be implicitly operating here; it also shows that there is no straightforward route that could take an agent from K to knowing a relevant set of probability claims. Examples like (*) can be easily multiplied; except perhaps for extremely simple cases, the richer probabilistic structure that typically superimposes on K can be crucial at the time of deciding what counts as an appropriate motivating reason.

Perhaps the idea is that we should just assume the existence of a prior objective probability distribution that provides the input to calculate the relevant conditional measures.⁸ Even though the authors never go down this path explicitly, we could seek to interpret their position as loosely based on this thought. However, according to this interpretation Hawthorne and Stanley should say that $P(p | K) = r$ can be treated as a legitimate reason for action only if the agent has knowledge of the objective priors on the basis of which suitable confirmation measures are obtained. But, as example (*) proposed above shows, in typical cases we cannot assume references to a hypothetical prior distribution to be implicit in standard discourse. Hence, at the very least, for the explanation of an action to make sense the agent would need to mention the relevant priors. Moreover, the lack of linguistic evidence can be taken to favor skepticism on the very existence of objective priors of the type required; postulating such measures might still turn out to be productive on theoretical grounds, but we are then left with the urgent task of discussing the details of such a theory and its relation with reasons as normally

⁸ For instance, Timothy Williamson has proposed an objective sort of Bayesianism, according to which we can identify a prior probability distribution that measures the intrinsic plausibility of hypotheses prior to investigation. See his *Knowledge and its Limits* (New York: Oxford, 2000), p. 211.

given by speakers – *precisely* because Hawthorne and Stanley place reasons at the center stage of their project.

To sum up, Hawthorne and Stanley make room for the intuition that probability claims can sometimes act as motivating reasons; in this case RKP demands that we know the corresponding probabilities, which are further construed as epistemic, in an objective way. But we have no indications as to whether such objective measures can be defined and are actually known. The authors try to circumvent this problem by focusing on our knowledge of non-probabilistic facts. However, just mentioning *K* is not enough – we still need explicit references to the appropriate measures, on pain of making a discourse about reasons unintelligible. So the problem remains. Moreover, it is not obvious that the linguistic evidence excludes a different, more subjective, interpretation of ordinary probabilistic discourse, in which case resorting to knowledge/belief-talk disregards crucial phenomenological aspects of the situation, as seen from the agent's point of view. In short, regardless of the merits of RKP for full beliefs, the attempt to squeeze all probability references (at the time of giving reasons) into RKP does not seem to be successful. The present reflections point to the fact that RKP cannot be the whole story on the link between knowledge and reasons for acting.