

A PUZZLE ABOUT NORMATIVITY

GIOVANNI ROLLA

Abstract. In this paper, I present a possible solution to the puzzle unveiled by Kornblith (2012) about the sources and the possibility of knowledge of epistemic norms. The puzzle is: if such norms cannot be discovered solely by reflection (“looking inwards”, thinking about first-order thoughts), and if there are correct ways of thinking and inferring, then such norms can only be discovered by investigating the world (“looking outwards”) —a counterintuitive conclusion. To avoid skepticism about normativity, I argue that we create normative correctness and discover normative demands by investigating the world and reflecting about our epistemic practices. This is done by an exposition of the method known as reflective equilibrium, which is defended against Kornblith’s thesis that the appeal to reflective equilibrium is doomed to failure because it implies reflection.

Keywords: Reflection; normativity; paradox; reflective equilibrium.

1. Introduction: a puzzle about normativity

In the book *On Reflection*, Kornblith (2012) argues against the idea that reflection is the source of normative demands, that is, that we can discover or create the correctness of normative claims by reflecting upon our first-order mental states (where a first-order mental state is a thought or experience that is not accompanied by the self-ascription of a thought or of an experience). This is a mixed point about the metaphysics and the epistemology of normative claims — whether they exist and how we can discover or create them —, and the following argument presents us with a puzzle about the sources of normativity (henceforth I will focus on epistemic norms, exceptions will be made explicitly):

- (1) Normative claims are not to be discovered only by investigating the world,
- (2) We are subject to normative claims,
- (3) ∴ Normative claims can only be discovered by reflection.

As Kornblith notes, it is an intuitive idea that normative demands are not like descriptions about physical objects — we cannot simply discover what the right way of thinking and acting is in the exact same way as we discover relations between physical objects and their attributes. As he puts it:

Normative demands do not seem to be part of the furniture of the world, something ‘out there’ which we might encounter in the way that we encounter tables and chairs, other people, or microphysical objects. The idea that there might be normative demands which have dominion over us is not simply some sort of causal fact, in the way that it is a simple causal fact about the world that the laws of physics have a kind of dominion over us and that we are subject to their constraints. We can, and sometimes do, violate normative requirements. We do not ever, nor is it possible to, violate the laws of physics. (Kornblith 2012, p.108)

This sets the background for (1). And if we want to preserve the idea that *there are* correct ways of reasoning and of acting, and that violating those ways makes one blameful — that is, if we want to maintain that we are subject to normative claims, (2), — the obvious conclusion is that we can only discover or create normative claims by reflecting upon our first-order thoughts, viz., (3). However, (3) cannot be the case, and Kornblith presents two compelling reasons why. First there is the regress argument: if a belief needs to be reflectively successful to have the correct normative status — if it needs to “pass the test” of reasons, so to speak — then the second-order beliefs called to reflect upon the first-order beliefs must themselves be reflectively successful, and this requires a further level of reflection. Thus, discussing Korsgaard’s proposal of Kantian inspiration for the problem of normativity of moral reasons, Kornblith writes:

When I reflect upon my reasons for believing [...] the reasons I have for that belief seem to be good ones. But surely this alone cannot constitute reasoning ‘all the way back’. I am aware that reflection can be too casual, and I am aware, as well, that reflection itself, just like first-order reasoning, can be biased and unreliable. So I wonder whether my reflective (second-order) assessment of my first-order reasons should be taken at face value. And this requires some sort of third-order assessment, which itself cannot be taken at face value. (Kornblith 2012, p.112)

To reason “all the way back” is simply to reach an unconditioned, indisputable result about the rationality and reliability of our beliefs, and it is far from clear whether finite beings like us can do such a thing — at best we can act as if we could reason all the way back, but this suggestion is not clear either. The idea here is also that the reliability and the rationality of our second-order beliefs, the ones involved in assessing the reliability and the rationality of our first order beliefs, can be called into question by the same kind of consideration that calls into question first-order beliefs: how can I be sure that the result of the reflective enquiry about what I should believe is itself reasonable? Couldn’t I have missed something? How can I know that my reflective assessment does not fail to identify the reasons to believe and whether those reasons are epistemically well grounded? Therefore, another level of reasoning is required, and so on.

Aside from the regress argument, the idea that reflection is a necessary condition for the normative status of a belief smuggles empirical presuppositions about what reflection can actually do. Reflective thinking is no more immune to familiar problems than first order beliefs are — it cannot, by itself, rule out wishful thinking, bias, unreliability and selective evidence. While it is quite intuitive that reflective assessment is a distinctive feature of human beings, hence of rational agents, empirical studies mentioned by Kornblith have shown that reflection does not guarantee a better sensitivity to reasons and an enhancement of reliability concerning the goal of truth. What is worse, when we try to assess our beliefs we often reason in favor of what we already believed, corroborating our beliefs by rationalizing them and misidentifying their sources (see Kornblith 2012, p.20–7, and specially Kornblith (1989), for a more detailed discussion). It is hard to imagine that under such conditions reflection could be of any gain in discovering or creating normative standards.

If both the conceptual and the empirical arguments set against (3) are sound, then we face the following problem: the argument from (1)–(3) can only be viewed as a RAA, and we have to choose between giving up (1) or (2) or even both. This is the reason why Kornblith's argument can be seen as puzzle or a paradox: on the one hand, if we reject (1), how are we to give any positive account of the idea that normative claims can be discovered just as we encounter “tables, chairs, other people” and so on? As he notes, it is indeed “mysterious how empirical investigation could ground normative claims” (Kornblith 2012, p.161). If we reject (2), on the other hand, we face skepticism and relativism, for if we are not subject to normative claims, it follows that there is neither *right* nor *wrong* way of thinking, reasoning and acting, and this is an uninviting conclusion — everything goes. Moreover, if the fact that (2) is false is the reason why the argument from (1)–(3) isn't solid, despite being valid, we should offer an additional explanation of why we *only appear* to apply normative concepts in appraising ourselves and others, epistemically and otherwise. For we clearly judge ourselves and others for behaving in certain ways, we reject some inferences and accept others, usually because some are good and some aren't. Why would we engage in such activities if normative claims have no power over us, or if their power is merely apparent? This is not simply a rhetorical question, it is a puzzle no easier to solve than the original one.

Kornblith himself is inclined to accept the denial of (1) as the correct diagnosis of his puzzle — saying nothing more than it is indeed a mysterious solution — and he does not consider the alternative of denying (2) (see Kornblith 2012, p.161). This is all too reasonable. However, there is another possibility, one unexamined so far, that aims to show that the distinction between looking outwards and “looking inwards” is not adequate, and since that the inference to (3) from (1) and (2) supposes this distinction, we might as well try and solve the puzzle on that basis.

2. Looking inwards?

When it comes to the topic of self-knowledge or introspection, there are available models that stress the similarities between knowledge of one's own mind and knowledge of the world in contrast with there being an insurmountable gap between those two kinds of knowledge. One such model has been influentially presented by Gareth Evans (1982), who argues that knowledge of oneself is much like knowledge of the world — *not* in accordance with the widely shared idea that “perception” of our mental states is analogous to perception of the world outside us, what would force us to postulate some kind of “inner sense” analogous to the outer senses, some capacity to scan the contents of our mind like our body is capable of scanning the world wherein we live. It is rather the other way around according to Evans: we come to know our mental contents and acquire self-knowledge by applying precisely the same procedures of determining whether the relevant *objective* state of affairs is the case.¹ As he famously wrote:

In making a self-ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outward . upon the world. If someone asks me ‘Do you think there is going to be a third world war?’, I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question ‘Will there be a third world war?’. I get myself in a position to answer the question whether I believe that *p* by putting into operation whatever procedure I have for answering the question whether *p*. (...) If a judging subject applies this procedure, then necessarily he will gain knowledge of one of his own mental states: even the most determined sceptic cannot find here a gap in which to insert his knife. (Evans 1982, p.225).

Whether the kind of case Evans has in mind is to be qualified as *knowledge* or as *justified true belief* or simply as *true belief*² — or even as *self-consciousness* without *self-knowledge* — will depend on how we construe these epistemic concepts, but this is not central to the problem we are facing here.³ Evans' model of introspection, which equates answering to the challenge of whether I believe that *p* with the challenge of whether *p*, represents an attack on the supposed adequacy of the distinction between looking outwards and “looking inwards”, for if one wants to discover the contents of one's own mind, there is no internal informational state to be “seen” by way of introspection. Clearly the phrase “looking inwards” can only be used metaphorically — and, for that matter, temerarily as well, for this usage of the phrase gives rise to the misleading idea that one can only know the contents of one's own mind by perceiving them as if they were internal objects. As the way we put it makes clear, the attack happens in a quite circumscribed domain — that of self-knowledge — but it may serve as an insight about reflection more generally. The suggestion is that to reflect about my belief that *p* is simply to apply the available procedures to decide

whether p , that is, the relevant conceptual and referential abilities, however inaccurate, biased and partial they are, that give rise to the belief that p (the same can be said about a chain of reasoning, understood as a sequence of first-order thoughts upon which we supposedly would reflect in order to assess their correctness). And if the first-order thought that p , which is the result of the application of those conceptual and referential abilities, cannot be the kind of thing that provides the correctness of normative claims, then second-order thought — understood as the exercise of the conceptual and referential abilities used in coming to believe that p — cannot either.

Clearly, this cannot be the whole picture about reflective appraisals, for sometimes reflection doesn't involve *self*-ascriptions of thoughts and experiences (we'll see more about that in the fourth section of this paper). Moreover, sometimes reflection, even when it's about one's own mental states, is about inferring consequences of our beliefs or stipulating alternative situations about what else *would be* the case if a believed proposition were the case. But what is central to the idea that reflection about first-order thoughts is necessary for their correctness is the assessment of the conceptual and referential abilities used in entertaining those thoughts — if the abilities were used in the right way —, and here the suggested model is applicable. This doesn't mean that every conceptual and referential abilities used in forming first-order thoughts are necessarily transparent to reflective assessment by the thinking subject. The conclusion is rather that the reliability and the correctness of a reflective processes is dependent upon those of first-order thoughts — what is compatible with the idea that sometimes reflective assessment cannot accurately identify the bases of belief formation or reenact the same procedures used in first-order thinking.

If the above suggestion is correct, then reflection is not something over and above first-order thought, for there seems to be no substantive epistemic difference between reflecting about my belief that p and believing that p : when it comes to certain propositional attitudes, there is no “looking inwards” that cannot be looking outwards. If this distinction is inaccurate, the argument from (1)–(3) is deemed invalid. Another construal of the argument would be of an enthymeme, exhibiting the following form:

- (1) Normative claims are not to be discovered only by investigating the world,
- (2) We are subject to normative claims,
- (2') If something is not to be discovered by investigating the world (looking outwards), it is to be discovered by reflection (“looking inwards”),
- (3) \therefore Normative claims can only be discovered by reflection.

I have argued that we have good reasons to regard (2') as false, and this is a seemingly satisfactory way to solve the initial puzzle, for the enthymematic argument depends upon a false premise (or, according to the alternative formulation, hinges

upon an illegitimate distinction). Before assessing this solution, we should turn now to some adjacent matters so far left behind.

3. Reflection and a priori

It is tempting to follow the discussion about the relation between self-knowledge and reflection and draw parallel conclusions about *a priori* statuses and reflection. The success of this approximation will depend on how we construe such concepts. Usually apriority is understood as a kind of justification or knowledge acquired independently of any experience. However, this characterization encompasses very different interpretations, and this gives rise to various problems. For instance, if we interpret apriority (of knowledge or justification) as the acceptance of a proposition p such that $\neg p$ is or implies a contradiction, clearly reflection isn't a priori. Take the example of someone who sees a tumbleweed on an empty road and catches himself wondering "this reminds me of a western movie". His attitude can be described as reflective, but the sentence "this *does not* remind me of a western movie" is neither a contradiction nor implies one. However, we do not need to interpret the a priori in this way. There are alternative interpretations, and although I will not defend such alternatives here, it is important to call attention to more mundane interpretations of apriority, for there is a possible point of convergence between reflection, self-knowledge and apriority.

Gilbert Harman (2003) has defended that strong accounts of apriority either are straightforwardly false and face skepticism or are reduced to the weaker accounts he calls *vanilla* conceptions, viz. (a) prior in relation to an enquiry and (b) biologically prior (see Harman 2003, p.24–5). In the first sense, something is a priori in relation to an enquiry if it is known or justifiably assumed before hand — e.g. the method applied in said enquiry and some propositions whose truth is supposed to be established beyond the scope of investigation, but whose status can be altered when new evidences are found or new conclusions are achieved. In the second sense, something is a priori if it is part of our constitution, that is, by being an innate disposition to accept certain appearances at face value. In both senses there is the possibility of error and correction, defeasibility. Consider a method M that was successful up until t in a domain D , but, given some variation in D at t_1 , M starts to display the wrong results. Someone who used M up until t in D had prior justification to do so, but not after the acknowledgment of its inadequacy. The same thing can be said about fundamental propositions assumed in an enquiry but subsequently proven to be false with the development of such enquiry or with parallel discoveries. Similarly, innate dispositions can be proven to be generally or locally unreliable, and it is the test of coherence between the results they offer and other beliefs that will force us to correct our innate abilities and their deliverances. Thus, if Harman's arguments are sound, a priori knowledge and justification do not guarantee infallibility.⁴

This result is important because it shows that a priori statuses are much like everyday knowledge and justification of contingent propositions, precisely in the sense they do not eliminate every possibility of error, just as the procedures of self-knowledge (considering some propositional attitudes) are the procedures of deciding whether objective states of affairs obtain, and as reflection is not epistemically superior than first-order thoughts. This is the point of convergence announced above, the enlightenment of these concepts and of what we can legitimately demand of them. In a word, demystification.⁵

4. Saving reflective equilibrium

We have seen one way to solve the puzzle about normativity, namely, rejecting the distinction between looking outwards and “looking inwards”, and that this is in tune with an account of a priori statuses that implies fallibility. However, even if we have solved the puzzle, the results we reached leave us short of giving a positive answer to the initial question: *what are the sources of normative claims?* Now I will answer this question by rejecting the dichotomy between the first premise and the conclusion of the puzzle, arguing that normative claims are not to be discovered exclusively by reflective thinking (as Kornblith has shown), but they are not to be discovered only by investigating the world either.

Let us remember that Kornblith sympathizes with the strategy of rejecting (1), but he acknowledges it is mysterious at best (2012, p.161). This comes as no surprise, for he also argues against the only available method capable of explaining how we could come to discover normative claims by investigating the world. This method is the one called *reflective equilibrium* by John Rawls (1971), and was influentially introduced by Nelson Goodman (1973) in the epistemological debate as how inductive and deductive inferences become epistemically justified, viz., by adjusting *both* particular inferences and general rules of enquiry through our cognitive practices. As Goodman puts it: “A rule is amended if it yields an inference we are unwilling to accept; an inference is rejected if it violates a rule we are unwilling to amend”. (Goodman 1973, p.64).

Kornblith’s argument is that, since reflection cannot be the source of normativity for this is both empirically implausible, given the real limitations of reflective thinking, and logically impossible, according to the regress argument, then reflective equilibrium is equally doomed if it implies reflection. He writes:

On Goodman’s account, then, the justification of both inductive and deductive rules is achieved by way of reflection, for we need to consider not only what inferences we actually make, and what inferential rules we actually endorse, but which changes we are willing to make when we find that there

is a conflict between our inferential practice and our beliefs about proper inference. Hence, reflection is put to epistemological work, in discovering what rules of inference are valid. (Kornblith 2012, p.122)

It is important to note that the case for reflective equilibrium is lost only if it implies a narrow notion of reflection, say, *self-referential reflective thinking*, where a subject describes exclusively his *own* inferential processes and, in cases of conflict, amends the relevant rules or reject the relevant particular inferences — how could he assure himself that the results achieved by self-referential reflective thinking are correct? This is clearly a herculean task — but not all reflective thinking is self-referential: we can, and indeed we do, reflect about the thoughts of others in order to assess the correctness of general rules of inference. There is a broad notion of reflection at play here, not out of tune with the way we use the phrase “to reflect upon”: it is not limited to self-ascription of mental states and processes, while it does accept the possibility of describing mental states and processes of others as a preliminary step to formulate and amend general rules of enquiry. If it is granted that we have a reliable intuition to recognize particular cases of correct inferences and paradigmatic instances of epistemic statuses (knowledge, justification, entitlement, what have you), to reflect in the third-person point of view means to describe such cases and formulate the general rules underlying them. When we reach the position of reflective equilibrium by reflecting (broadly conceived) we are not trying to submit our *own* inferences to the test of reasons, what would makes us prey to the regress argument, we are discovering processes that are more or less established in the practices of the community. By doing so, we are being part of the community, and our own enquiries and the way we act based upon such enquiries are up to evaluation by others. In construing reflective equilibrium this way we avoid the empirical argument as well: we minimize the risks of being recurrent victims of wishful thinking and are confronted if our beliefs are biased or unreliable and so on.

Of course, even if Goodman’s original idea is that we have to reflectively assess, in the narrow sense, the results of our enquiries and choose what to believe in, we do not need to follow him about this, no matter how tempting it is to do so — take for example the description of reflective equilibrium by Gilbert Harman:⁶

In deciding what to believe or what to do, you have to start where you are with your current beliefs and methods of reasoning. These beliefs and methods have a privileged status. You are justified in continuing to accept them in the absence of a serious specific challenge to them, where the challenge will typically involve some sort of conflict in your overall view. Conflict is to be resolved by making conservative modifications in your overall view that make your view more coherent in certain ways. (Harman 2003, p.25–6).

The problem begins with the vocabulary: *to decide what to believe* seems to be a matter of reflective appraisal in the narrow sense. It seems to imply that one has to

be able to recognize a tension in one's practices and freely decide whether a general rule is to be amended or a particular inference is to be rejected. What is especially problematic in this construal of reflective equilibrium is the supposition that we are able to have some kind of voluntary control over our propositional attitudes, what is obviously false for a large class of cases, namely, perceptual beliefs, beliefs acquired through memory and sometimes even through testimony. But the kind of reflection involved in reflective equilibrium needs not to imply the act of deciding so conceived. And what exactly would force us to refine our general rules of enquiry and reject particular inferences, if not a voluntary act to do so? As Goodman calls it, in cases of conflict we are guided by our *unwillingness* to amend a rule or accept an inference, that is, a disposition to conserve our stocked body of information and knowledge rather than radically revising it at every turn, what is a matter of practice, not of postulating a rule and conscientiously testing and comparing the results. This covers the problem of amending our inferences without implying reflection (narrowly conceived), but how can we discover normative demands by investigating the world as well as reflecting (broadly conceived) about our inferences?

To answer to that, let us begin by noting that the characterization of reflective equilibrium as the source of epistemic justification of inductive and deductive rules can be understood as the source of normative correctness. The reason for this is that a particular inference is *correct* if it conforms to the relevant rules we are willing to preserve and a rule is *correct* if it covers a sufficiently large class of cases we are willing to accept. And if precisely what rules are to be amended and what inferences are to be rejected depends on our disposition to accept the deliverances of our enquiries, this clearly involves testing, revising and being sensitive to new information. We have to look outwards. That is not to say that we discover what makes a normative claim correct in the *exact* same way we discover what makes the description of a physical object true, for the difference lies in the fact that we do not discover truth-makers of normative claims, we *entrench* — to borrow from Goodman — normative correctness through our cognitive practices, adjusting them diachronically. And only when we have a substantive record we are able to identify paradigmatic cases, describe our cognitive procedures and formulate normative claims — such as “it is correct to believe that p iff one has evidence on the kind E ” — and if this process results in a *discovery*, it is heavily dependent upon the development and history of our practices. Even at this later stage, in which philosophers, cognitive psychologists and neuroscientists are put to work, self-referential reflective thinking is not necessary, for the description of our practices as a preliminary step to express a precise formulation of what is the correct way of reasoning does not need to (and often *does not*) involve any self-ascription of present or past experiences and thoughts.

But what about the amendments we make? Surely, if we look back at the epistemic restrictions at play in non-deductive inferences for instance — diversity and size of

the sample, relevant similarities (in cases of analogical reasoning), simplicity of explanation (in cases of abductive reasoning), regularity in observed cases (in cases of projective reasoning) and so on — we cannot but be in awe about the complexity of such refinements. It is not at all puzzling to suppose we have encountered particular problems and particular revisions have been made, and what we now have is the conjunction of those amendments after they have been entrenched and refined. But even if the possibility of addressing particular problems as they appear and become more present is regarded as too far-fetched, the reflective equilibrium made by specialists is able to arrive at such restrictions by identifying and describing paradigmatic cases of successful inferences without self-referential reflective thinking.

Kornblith himself has somehow anticipated the strategy of solving the puzzle by arguing that normative claims are to be discovered both by reflective thinking and investigating the world, but quickly dismisses it. He writes:

Normative truths might thus be identified with the stable results of reflective normative inquiry, not by standing back from one's beliefs and desires at an instant and reflecting on normative questions without further empirical input, but by engaging in normative reflection with the aid of further empirical input. The normative results one reaches must thus withstand not only reflection on one's current state, but they must remain stable even in the face of further engagement with the world. (Kornblith 2012, p.118)

In this exert Kornblith clearly has in mind the narrow conception of normativity rejected above — but his dismissal of the proposal is surprisingly based on the possibility of misleading empirical enquiry (which might include, but isn't limited to the feedback of one's epistemic community and so on):

Even this more stringent requirement will not guarantee convergence of any kind, let alone convergence on normative truths. Empirical inquiry may turn up misleading evidence of various sorts, and so it need not inform, but may instead distort, the results it produces. The idea that empirical investigation is likely to correct mistakes which might otherwise result is overly optimistic [...] The idea that such an investigation must inevitably correct any errors is clearly just mistaken. (Kornblith 2012, p.118)

Set aside the fact that an essential part of his argument against (3) above is empirically informed, to argue that empirical studies cannot help our cause for normativity by appealing merely to the fact that empirical studies *might be mistaken* is an overly *pessimistic* conception of epistemic statuses. We justifiably believe all sorts of things in our everyday life without excluding every possibility of mistake — and if Kornblith is flirting with infallibilism, the thesis that we can only know or justifiably believe a proposition if we hold an infallible epistemic relation to it, then skeptical intuitions are underlying his argument. Moreover, if the empirical sciences are not of any guidance, *what else* could be?

Finally, there is another objection to consider: if we describe our epistemic progress this way we seem to open too much room for *lucky* guesses, for how are we to be certain that the proper amendments have been made and not just *ad hoc* alterations to cover particular results? Sure, one can answer, if we did *only* that, we would soon face more counterintuitive results without being able to deal with new cases, this way a radical revision of our practices would be inevitable. That is so because to cover new, unpredicted (and so far unconceived) cases preserving the core general rules available is a valuable epistemic end — if our procedures repeatedly fail at this goal we would have to start from the scratch more times than what is practically recommended.

This answer is correct, but it misses the point of the objection: how can we *know* we have excluded luck? For we cannot appeal to our epistemic practices and their deliverances to show that we have successfully avoided lucky results: this would be blatantly circular. Moreover, if knowledge is incompatible with luck, and if our practices do not offer the guarantees that we have excluded any lucky reasoning, they leave us inevitably short of knowledge. What good are they then?

First of all, it goes without saying that a complete defense of reflective equilibrium has to be able to distinguish an argument or inference that is *viciously* circular from an argument or inference that is *virtuously* circular. While it is clear that reflective equilibrium is indeed the best available answer to the question of how we came to create normative correctness and discover normative demands, it still has to be shown that the circularity of appealing to our own practices in order to show how successful they are is virtuous rather than vicious, but to argue thus is beyond the present scope of investigation.

Now, the core problem presented by the objection requires us to consider only what luck is, especially concerning knowledge and other epistemic statuses (such as justification and entitlement). The most complete account of epistemic *luck* is due to Pritchard ((2005), chapters 5 and 6), and it offers us a good guide here. There is indeed a clear sense in which knowledge is incompatible with luck, namely: a true belief does not amount to knowledge if it could very easily be false.⁷ Such are the famous Gettier cases, in which a subject justifiably believes in a true proposition that would be false in similar situations — and, in these situations, he would retain his belief in the same proposition (hence believing falsely).⁸ The idea that a belief could very easily be false is to be understood through the modal vicinity of the relevant event, and this vicinity has to be organized in certain way. In Gettier cases, there are a lot of possible scenarios wherein the subject believes falsely by applying the same *method* of belief acquisition. That the method is fixed is important, because we can imagine alternative scenarios where the subject arrives at the same (true) belief by reliable ways, diminishing the possibility of believing falsely. Such considerations may help us as follows: if our procedures are plagued by epistemic luck, then they deliver

the right results but could very easily deliver the wrong results. If the modal vicinity here is to be organized by the procedure (method) under scrutiny, then there are nearby possibilities in which our procedures would give us the wrong results. Assuming that nearby possibilities are probable events, it is clear that we would sooner or later face such counterintuitive results and then we would be forced to make amendments, which would eventually minimize the effects of epistemic luck. However, if the possibility of delivering the wrong results is not modally close (alternatively, is not probable), thus being far-fetched, then there is no need to worry at all: for one thing, our methods surely are not infallible (we can say we are *a priori* justified, in the senses (a) and (b) seen above, to infer as we do) — on the contrary, their fallibility is a condition for their progress towards the correction of epistemic norms. Secondly, if the possibility of mistake is remote, by definition it is not a matter of luck that we get the right results.

5. Concluding remarks

We started with a puzzle about the source of normative demands and focused our argumentation on epistemic normativity. Kornblith's argument is correct up to some point: if one accepts that we cannot create or discover normative claims by investigating the world, then one has to accept that we cannot create or discover normative claims only by "looking inside" for, as we've seen while discussing Evans' model of self-knowledge, there is no substantive epistemic difference between reflection from the first-person point of view and first-order thinking — the reliability of the former depends on the later. Secondly, his diagnosis of the puzzle is correct as well: the only strategy of solving the puzzle that explains what are the sources of normativity — aside from embracing skepticism and relativism — is to accept that we discover norms by investigating the world, at least partially. Kornblith interprets the process known as reflective equilibrium, the only available method capable of explaining how we can discover norms by looking outside, as implying a narrow conception of reflection, and I have argued that this is not the case. The variation of reflective equilibrium defended here explains how we are able to arrive at normative claims by reflecting *and* investigating the world. Very little has been said about *moral norms*, and as of now we can hope to develop similar arguments and achieve similar results on this subject.

References

Bar-On, D. 2004. *Speaking My Mind: Expression and Self-Knowledge*. Oxford: Oxford University Press.

Principia 18(3): 323–336 (2014).

- Carruthers, P. 2011. *The Opacity of Mind*. Oxford: Oxford University Press.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.
- Gettier, E. 1963. Is Justified True Belief Knowledge? *Analysis* 23(6): 121–3.
- Goodman, N. 1973. *Fact, Fiction and Forecast*. Indianapolis: Bobbs-Merrill.
- Harman, G. 2003. The Future of the A Priori. *Philosophy in America at the Turn of the Century, Centennial Supplement to Journal of Philosophical Research*: 23–34.
- Henderson, D.; Horgan, T. 2011. *The Epistemological Spectrum*. Oxford: Oxford University Press.
- Kornblith, H. 1989. Introspection and Misdirection. *Australasian Journal of Philosophy* 67: 410–22.
- . 2012. *On Reflection*. Oxford: Oxford University Press.
- Kripke, S. 1980. *Naming and Necessity*. Oxford: Basil Blackwell.
- Medeiros, E. V. 2013. *Teorias da Introspecção e psicologia moral*. Retrieved from [https://www.academia.edu/4069146/Teorias da Introspeccao_e_psicologia_moral](https://www.academia.edu/4069146/Teorias_da_Introspeccao_e_psicologia_moral)
- Pritchard, D. 2005. *Epistemic Luck*. Oxford: Clarendon Press.
- Rawls, J. 1971. *A Theory of Justice*. Cambridge: Harvard University Press.

GIOVANNI ROLLA

Universidade Federal do Rio Grande do Sul
rollagiovanni@gmail.com

Resumo. Neste artigo, eu apresento uma possível solução ao problema criado por Kornblith (2012) sobre a origem e a possibilidade de conhecimento de normas epistêmicas, seja ele: se tais normas não podem ser descobertas por reflexão apenas (“olhando para dentro”, pensando sobre pensamentos de primeira ordem), e se há modos corretos de pensar e de inferir, então tais normas apenas podem ser descobertas pela investigação do mundo (olhando para fora) — uma conclusão claramente contraintuitiva. Para evitar o ceticismo sobre a normatividade, eu argumento que nós criamos e descobrimos demandas epistêmicas ao investigarmos o mundo e refletirmos sobre as nossas práticas epistêmicas. Isso é feito através de uma exposição do método conhecido como equilíbrio reflexivo, que é defendido da tese de Kornblith de que o apelo ao equilíbrio reflexivo está fadado ao fracasso porque implica reflexão.

Palavras-chave: Reflexão; normatividade; paradoxo; equilíbrio reflexivo.

Notes

¹ Evans incessantly reminds us that this cannot be the whole picture about *self*-knowledge (see, for instance, Evans 1982, p.226). It is also needed that a subject who entertains a thought be capable of envisaging what would be for other persons and for himself at different times to entertain the same thought (a similar constraint is applied to self-ascription of perceptual experiences as sources of self-knowledge). This is the requirement that Evans calls *General Constraint*, and it enables the personal pronoun in ‘I think that *p*’ to refer to the subject of ascription, hence being more than a meaningless attachment to *p*.

² Evans actually uses a primitive notion of knowledge as information so the standard tripartite analysis of knowledge may not apply (see Evans 1982, p.122–3), though it is an easy prey to Gettier cases.

³ The limited kind of cases this model is able to deal with can be a compromising feature as well, but this is not a problem for our present argument, for we are not interested in an account of self-knowledge, and *a fortiori*, not interested in a complete account of self-knowledge either. Evans himself thought that his conception of self-knowledge (acquired through self-ascription of thoughts and self-ascription of perceptual experiences) would serve for other propositional attitudes, but the literature has shown otherwise — see Bar-On (2004) and Carruthers (2011). Other propositional attitudes that don't fit Evans' model of self-knowledge do not bear on our matter. I owe this observation to Eduardo Vicentini de Medeiros (Cf. Medeiros E. V. 2013, p.51–7).

⁴ This leads us to the conclusion that there is more to the contingent a priori than what we've already learnt from Kripke (1980).

⁵ Other accounts of the role of apriority have surfaced recently. Henderson, D. & Horgan, T. (2011, chapter 2) argue that, whether or not there is a kind of knowledge or justification that passes the strongest interpretations of apriority, much of the post-Gettier epistemological enquiry is made a priori but with *empirical infiltration*. This means that whenever epistemologists propose a definition of an epistemic concept, new information provided by thought-experiments and counterexamples, noise removal concerning language use, among other phenomena, are weighted in the assessment of the proposed definition. The options usually are refining it, rejecting it or accusing the new information (mainly in case of counterexamples) of being underdescribed, thus preserving the definition. When the result is successful, it offers a defeasible a priori justification for accepting the proposed definition.

⁶ Harman defends his version of reflective equilibrium by saying that we are *a priori* justified to infer as we do, according to the definitions (a) and (b) of apriority mentioned above — but I don't take sides here if the method of reflective equilibrium is more properly characterized as *a priori* or *a posteriori*.

⁷ Pritchard argues that there is another variation of luck that is ineliminable, namely, reflective luck — that from our *reflective point* of view we cannot exclude certain hypothesis that would render our everyday beliefs false — even if our everyday beliefs do exclude luck in the first sense (Pritchard 2005, p.173–8). Pritchard concludes that we cannot rule out reflective epistemic luck with an argument based on the Underdetermination Principle (for any subject S, propositions *p*, *q*): If S's evidence for believing *p* does not favor *p* over an incompatible hypothesis *q*, then S is not justified in believing that *p*. Since we do not have (reflectively accessible) evidences to believe that *p* that rule out an incompatible skeptical hypothesis, such as the brain in a vat hypothesis, we are not internally justified to believe that *p*. This problem emerges because the elimination of reflective luck is a matter of possession of reflectively accessible justification, but it is not clear whether we should accept this.

⁸ Cf. Gettier (1963) for the original cases — but countless other cases have appeared since then.