

Comment on Paul Boghossian, “What is inference”

Crispin Wright

© Springer Science+Business Media B.V. 2013

Abstract This is a response to Paul Boghossian’s paper: What is inference? (doi: [10.1007/s11098-012-9903-x](https://doi.org/10.1007/s11098-012-9903-x)). The paper and the abstract originate from a symposium at the Pacific Division Meeting of the APA in San Diego in April 2011. John Broome was a co-commentator.

The epistemology of logic presents a trifecta of major, interrelated problems: whether we know, and if so how we know, that our basic principles of inference are sound; how and when it is possible to use logic to increase our knowledge; and the nature of inference itself. In other work, Paul Boghossian has shed much light on the first two of these. In his present contribution he tackles some fundamental issues regarding the third and has, as usual, done much to illuminate the challenges posed.

My remarks will fall into three parts:

- On Boghossian’s set-up of the issue: on his focussing on the problem of inference as one of characterising a certain kind of *belief-formation*.
- On inference as a form of rule-following and the problem of the associated vicious regress. Can thinking of the processes involved as sub-personal help? Boghossian thinks not—at least not if our concern is with “person-level reasoning, reasoning as a mental action that a person performs, in which he is either aware, or can become aware, of why he is moving from some beliefs to others.”¹ I think this assessment is correct, provided any account of inference

¹ Boghossian, this volume.

C. Wright (✉)
New York University, New York, NY, USA
e-mail: cw43@nyu.edu

C. Wright
The Northern Institute of Philosophy, University of Aberdeen, Aberdeen, Scotland, UK

- needs to be constrained by Boghossian's Taking Condition. But I think the right conclusion is that the Taking Condition should be dropped.
- On a certain comparison, which I think illuminates *why* the Taking Condition should be dropped, between inference and rational action.

1 Inference as a kind of judgement

So first, on what may appear an awkwardness in the way Boghossian frames the discussion. In his initial toy example, the inference of

(3) The streets are wet

from the premises, respectively remembered and independently known, that

(1) It rained last night,

and

(2) If it rained last night, then the streets are wet,

is identified as involving a judgement of a proposition—that the streets are wet—on the basis of an acceptance of others—the two premises. And this is subsequently crystallised in the modified version of Frege's characterisation that Boghossian calls:

(*Inferring*) S's inferring from p to q is for S to judge that q because S takes the (presumed truth of) p to provide support for q.

Now, this may seem to mistake the target. Shouldn't we distinguish inference in general from *coming to a conclusion* in particular? One may infer from a supposition, or from a proposition one takes to be false (for instance, in the course of looking, perhaps unsuccessfully, for consequences that bring out its falsity.) Inferring is a movement of thought between propositions which may, in special circumstances, result in the thinker coming to judge the proposition inferred to be true. But no particular attitude to that proposition is implicit in inference itself, in particular not judgement of its truth. What Boghossian is talking about is coming to believe a proposition *on the basis of* inference. That is not inference per se. Inference is rather, one would suppose, a proper ingredient in it.

There is another point. Often, even when inference does result in a judgement of the truth of a conclusion, there remain no undischarged premises for that conclusion (when, for instance, one reasons by *reductio ad absurdum*, or conditional proof). In such cases, the conclusion is the product of inference. But one is not judging it true because one takes some premises to provide support for it. One has discharged the premises. So in such a case inference, at least in the sense characterised by (*Inferring*), seems not to be involved.

These points might be taken to suggest that Boghossian is starting in the wrong place: that, whatever his declared interest,² what he *ought* to be interested in is whatever constitutes the apparent *common factor* in the three cases, (i) of judging a conclusion true because one has inferred it from propositions one judges to be true, or (ii) reasoning from a supposition to its consequences, or (iii) of reasoning to the judgement of a conclusion in a way that discharges all other premises, judged true or no. But I think any such quest for a common factor is doomed. What is the alleged common factor? When inference is taken to be a process in common to all those three types of case, but something to which the judgements in the first, and the suppositions in the second are *additional*, and which provides the basis for the judgement in the third and so is in some sense concluded in advance of that judgement, then it comes to seem utterly elusive in what it could consist. There may be phenomenological commonalities in particular examples—the occurrence of an orderly series of thoughts, perhaps, with something of a sense of sequential discipline of some kind—but that hardly seems essential.

But how then to proceed? I think that, the foregoing misgivings notwithstanding, (*Inferring*) is a not unreasonable first shot—at least, it does not miss through lack of generality.³ But we need to do two things to accommodate the cases of suppositional inference, and premise-discharging inference to the model. The first is to generalise the formulation so as to allow inference to conclude in other kinds of attitude—more specifically, to allow other kinds of *acceptance* besides judgement, or belief—to eventuate from inference. Supposition, in particular, may—I think should—be viewed as a kind of conditional, or context-confined, acceptance. If one begins from suppositions and then, without discharging them, reasons validly to a certain consequence, then what is mandated is an acceptance of that consequence in the same context of supposition. So I propose, with acceptance so understood, that (*Inferring*) be generalised to something along these lines:

(*Inferring*)* S’s inferring from p to q is for S to accept q because S takes the (accepted truth of) p to provide (contextual) support for (acceptance of) q.

What about premise-discharging inference? Here the needed reflection is that it’s a mistake—or an artefact of terminology—to regard such inferences as, so to speak, eliminating anything in the role of a premise. One’s judgement of the truth of a conditional, for example, established by a step of conditional proof, rests precisely on the anterior judgement that the consequent has been validly derived from a pool of assumptions including its antecedent. The antecedent is discharged, but that anterior judgement is not. And it is arguably the latter from which the conditional is actually inferred.⁴

² Boghossian writes, “By ‘inference’ I mean reasoning with beliefs. Specifically, I mean the sort of ‘reasoned change in view’ that Harman (1986) discusses, in which you start off with some beliefs and then, after a process of reasoning, end up either adding some new beliefs, or giving up some old beliefs, or both.”—Boghossian, this volume.

³ I mean that it is not disposed of by the concerns about suppositional and premise-discharging inference. I do it mean that I think it is correct.

⁴ So much is explicit when reasoning by conditional proof is formulated in the notation of sequent calculus. Analogously for *reductio*.

So I don't think (*Inferring*) does in fact, in any serious way, undergeneralise the target phenomenon.

Inference, as characterised by (*Inferring*) or (*Inferring*)*, is subject to the Taking Condition that Boghossian formulates as follows:

Inferring necessarily involves the thinker taking his premises to support his conclusion and drawing his conclusion because of that fact

It is this, claimed necessary condition⁵ for a movement of thought to count as inference on which I want to focus.

2 The Regress

The regress adverted to by Boghossian derives from the combination of the thoughts

- (A) that all inference is a form of rule-following, and
- (B) that all rule-following involves inference,

—specifically, that all rule-following involves inferential transitions from states of mind (for now, personal states of mind: explicit acceptances and intentions) that carry the content of the rule(s) the thinker is following. Let's spell the regress out in fully personal terms.

Suppose I set myself in a particular context to follow rule R—say, Boghossian's email rule. In order to follow rule R on a specific occasion, I must, by (B), move inferentially—this is inference₁—from, *inter alia*, a state of mind that carries the content of R; and to make that inference₁, I must, by (A), follow an appropriate rule of inference, and so I must, by (B), make an inference₂ from a state that carries the content of *that* rule of inference. This inference₂ must, by (A), involve an episode of rule-following, and so must involve, by (B), an inference₃ from a state that carries the content of the relevant rule(s)... And this series of inferences is regressive because each of them has a distinct conclusion: my arrival at a mandate for its predecessor. Thus (A) and (B) combine to turn ordinary inference into a 'supertask'.

Now, clearly, bracketing the issue of vicious regress, this is already a parody of anything recognisable as ordinary rule-following. There is much too much thinking involved! Our actual following of rules, including inference if that is indeed a form of rule-following, falls into no such orgy of personal-level information processing and explicit reasoning. What the Regress shows is that it *could* not do so. But it should be obvious that there is no prospect of assistance with the bind merely in the suggestion that, in order to do better justice to the phenomenology, we should recast the model in terms of sub-personal processes and states. If sub-personal rule following involves sub-personal inference and sub-personal inference is itself a

⁵ Boghossian doesn't discuss the question of sufficiency, and nor will I here. But here is a question: suppose I accept certain propositions and then take it *on testimony* that they support a certain conclusion and proceed to accept it on that basis. Is the latter acceptance inferential?—I haven't, after all, actually carried out the inference myself.

form of sub-personal rule-following, then the threat of regress remains. Sub-personal systems—or anyway, finite populations of them—are no better at supertasks than persons are. The Regress will not be checked by retaining the *structure* of the processes of rule-following postulated in the personal case and then merely insisting that the model must sometimes be realised sub-personally rather than personally. If going sub-personal can help, it will be because it somehow allows us to refashion the structure of the processing.

But how might it do that? Let’s suppose that Boghossian is right that anything worthy of the term “inference” must be subject to something in the ball-park of his Taking Condition: again,

Inferring necessarily involves the thinker taking his premises to support his conclusion and drawing his conclusion because of that fact.

Then all inference has to involve if not a personal, then, at the least, some form of sub-personal ‘taking that’ concerning the obtaining of an appropriate support relation between premises and conclusion—(it won’t matter how exactly we characterise that support relation). Now I confess to seeing no alternative interpretation of this “taking that” than to say that it requires an *information-bearing state*: more specifically,—if we are construing matters sub-personally,—that the sub-personal system involves a state that, as we may say, *registers the obtaining of* an appropriate support relation or—to allow us to speak to the case where the inference is bad—that it *represents* such a relation as obtaining. (I’ll keep the term, “registration”, but you should avoid hearing it factively.) More austere, just think of the registration state as carrying a content that somehow licenses—gives a green light to—the relevant transition. The Taking Condition, then, is saying that all inference, personal or sub-personal, involves such a registration state.

It merits remark that, in order to do justice to the “because” in the Taking Condition,—“... drawing his conclusion because of that fact “,—we will have to add that this registration state somehow *controls* S’s movement in thought from judgement of the truth of his premises to judgement of the truth of the conclusion. And for the familiar reasons, ‘control’ here cannot just mean: *causes*—the causation has to be of the right, ‘non-deviant’ kind. So that vexed distinction is certainly going to remain on the agenda.⁶

But our concern is with the regress. And the threat re-emerges as soon as we ask the question: what *kind* of content is to be carried by the registration state?

We confront a dilemma. Suppose the content of the registration state is *general*: for instance, that any transition of the appropriate kind is licensed when the system is in a state of acceptance of the relevant kind of premises. Then it seems that we will have to understand the control exerted by the registration state on the specific movement in question as mediated by an ‘appreciation’ that the latter comes within the ambit of the former—so as an instance, in effect, of the inference: transitions of such and such a kind are mandated; this is a transition of such and such a kind; so it is mandated; so...[movement]. And now, fatally, we have represented the original inference as involving another; regress ensues.

⁶ As Boghossian is aware; see his n. 3.

If, on the other hand, the registration state is conceived as carrying a content *specific to the inferential transition* that it controls, the question we need to focus on is: how does the system get into the registration state in the first place? Our inferential capabilities are of course finite in the sense that they are bounded by limitations of time, concentration, intelligence and so on. But they are open-ended in the way in which our linguistic competences are open-ended: to have a specific capacity of inference is to be able to handle—make, and ratify—no end of inferences of the appropriate kind. If each of these performances requires the control of an appropriate *specific* registration state to count as inference, then we seem forced—for reasons exactly analogous to those that apply when we are trying to provide a theoretical representation of the ability to parse no end of novel utterances in our mother tongue—to conceive of these specific states as accessed by inference from general information states that encode our overall inferential program. But then we again have moved to think of any inference as underwritten by another—an inference to the relevant specific registration state—and regress ensues once again.

Provided, then, that we continue to insist that the Taking Condition remains in force as a condition on all inferring, properly so termed, and it is interpreted as requiring, at a minimum, that all inference is controlled by a suitable content-bearing registration state, it remains uncertain that any coherent—regress-free—model can be given of what inferring actually is. I say the modest “uncertain” rather than something more factive because of the slack in the notion of control: if there is a way that content-bearing states may intelligibly exert control on a system in line with their content, so to speak, without that system being required to *process* the content, then we may escape with the Taking Condition intact. But notice that the movements involved in the exercise of this putative form of control—whether they run from a registration state encoding a general content to the actual inferential movement, or from the states encoding the general inferential program to registration states of specific content that control specific inferential movements—will now need to be conceived as a kind of rational transition which is not subject to the Taking Condition: something the system quite properly does without any vestigial ‘taking it that’ the operations in question are in good standing. And if we are prepared—indeed forced—to countenance such operations at some point in any coherent model of inference, what motive is there to reject the idea that inference itself can be such an operation?

None is evident. But that is to say that, after all, there is no sound motive for insisting on the Taking Condition in the first place. We can, if we like, reserve the term, ‘inference’, for operations that comply with it, in some form. But we know that, at some level, our thinking must involve movements from and to information-bearing states that are, or so we hope, appropriate to the information those states carry but are not policed by any form of ‘taking it that’ they are so appropriate. And now there is no good reason to deny that some of the moves we call ‘inference’—perhaps those of Boghossian’s Hide-and-Seek playing children—are of this kind.

In summary: I believe that, its plausibility notwithstanding, we must drop the Taking Condition on inference—at least so long as it is understood to require that inference, properly so termed, is invariably under the control of an information state that somehow registers a mandate for the transition concerned. And if the Taking

Condition, understood as requiring such informational control, is a necessary ingredient in rule-following in general, properly so termed, we must drop the idea that inference is, everywhere and essentially, a kind of rule-following. That, in outline, is the solution to the problem of the Regress.⁷

3 Inference as rational action

But now we seem to be back at Square One. How, without the Taking Condition, are we to respond to the central question: that of distinguishing inference, properly so regarded, from a mere sequence of thoughts, even when later elements are viewed as caused by earlier—how are we to distinguish inference from mere mental ‘jogging’, as John Broome puts it? And how are we to save the idea of inference as something that *we do*, as a kind of mental action, rather than a series of mental events?

There is a proposal which Boghossian canvasses, and rejects, in the second section of his remarks, which, while I agree it misfires, points in what I think is the right direction. This was that a transition from a set of beliefs to another belief is inferential when the thinker regards the former as justifying the latter, as providing reasons for the latter. Boghossian termed this “the Counterfactual proposal”. The relevant counterfactual associated with an inference is glossed as something like: were the thinker to be asked why he accepts the putative conclusion, he would offer the premises as his reasons. Boghossian’s objection to this is that it represents the question, whether a particular transition is inferential, as response-dependent—as turning on facts about the thinker’s judgements—and that no such account can be correct if, as it seems, inference is a causal notion: if a necessary condition for inference is that the thinker’s acceptance of his premises actually produce his acceptance of the conclusion. For causality surely cannot be a response-dependent matter.

This is a strong objection (to a flimsy proposal). But consider instead the proposal, not that the status of the transition as inferential depends on the thinker’s judgments about his reasons, but that it depends on *what his reasons are*. We want his acceptance of the premises to supply his *actual* reasons for accepting the conclusion. What he is prepared to judge about his reasons is one, very good piece of evidence—in a wide class of cases, decisive evidence—about that. But it need not be part of the proposal that it is constitutive.

Call this the Simple Proposal. It says that a thinker infers q from $p_1 \cdots p_n$ when he accepts each of $p_1 \cdots p_n$, moves to accept q , and does so for the reason that he accepts $p_1 \cdots p_n$.

Here it is important that this is to be a complete explanation of the thinker’s acceptance of q ; it will not count as an inference if the transition additionally includes desires among its reasons—as in Pascal’s Wager—nor as an inference exactly from $p_1 \cdots p_n$ if its reasons include additional acceptances.

⁷ This matter is discussed further in Part I of my Replies to Commentators in Coliva (2012).

What is needed, then, is an account of, or at least some insight into, what it is for certain intentional states of a thinker to be his actual reasons for his transition to another intentional state. And the account we need must avoid returning us to the Taking Condition—it must avoid the idea that what we may term the actual-reason relation needs to be underwritten by an intentional state, personal or sub-personal, whose content is that the transition in question is in good standing, or that the initial states, or their content, support the end state. We need to avoid committing to the notion that doing something for certain reasons must involve a state that somehow registers those reasons *as* reasons for what one does. But if no such state need be involved, what does need to be involved?

Well, I think that exactly there is where the real problem lies: to understand what it is to act on certain specific reasons and no others. Inference is, at bottom, just a special case of that and I have no further account or analysis of it to offer here. But I think it may be illuminating briefly to consider the case of action in general, in three respects.

First, an action proper is distinguished from a ‘mere’ (as we like to say) bodily movement exactly in that it is done for certain reasons. The notion that inference is something that we do is immediately safeguarded by the Simple proposal, as a special case of the platitude that *an action* is something that we do.

Second, suppose a subject acts on a belief B and a desire D in such a way that B and D provide the reasons for her action. She wants a beer, maybe, and believing that there is a six-pack in the ‘fridge in the kitchen, heads off in that direction. To say that B and D provide the reasons for her action is to say that the subject’s performance is actually explained by her being in states B and D and that, so explained, it is to that extent rational. But it is not necessary, and in many cases would be incorrect, to advert in addition to her appreciation that these states do confer on her good reason for her journey to the kitchen. Action for reasons is one thing; awareness of those reasons as reasons, and especially as sufficient reasons, is something extra that need not be involved. A child, or a smart dog, may have reasons for action, and act on them, with no concept of reason in general, still less of its having reasons for what it does. There is simply no requirement for rational action in general of any analogue of the Taking Condition.

But it is the third point that is, I think, most illuminating. It is possible that our beer-seeking subject have had exactly that belief and desire, and had no other relevant beliefs and desires, and yet failed to head for the kitchen. That would be an example where the subject violated what we may term the norm of *Anti-Akrasia* (A-A):

Provided there is no overriding reason not to, do that which you believe will satisfy your desires.

It seems harmless and correct to say that normal agency is constrained to comply with the norm of A-A. But we should be very circumspect about the idea that the manner of this constraint is that, in ordinary action, we *follow* A-A. Whatever the subjection of rational action to A-A involves, what it does *not* involve is that we need to advert to some kind of *acceptance* of the norm in giving a full rational explanation of what a normal agent does, or in articulating the full practical

syllogism that provides that explanation. There is no analogue of the Taking Condition for A-A. A full rational explanation of an action need cite no more than the agent's relevant specific beliefs and desires; it does not need to add any mention of an intentional state—an intention itself, or an acceptance of a rule—that encodes A-A.

Now, Boghossian seems to me right that anything worth describing as rule-following had better involve *some* form of acceptance of the rule—some kind of registration of the content of the rule, and some kind of 'taking it' that *this* is what the rule requires—featuring essentially in the best explanation of the subject's performance. So A-A gives us a prototype of a principle that *in some sense* constrains—rather than merely describes—our practice but not on a rule-following model of that constraint. There is no such registration—no tacit intention to follow A-A, or whatever—in normal, non-akratic action. Rather, it seems we should say, it is *constitutive* of rational action that one behaves in ways that can be explained without citing any such intention, just by adverting to the relevant beliefs and desires. A-A is upheld by the pattern in the relevant kind of rational explanation; it is not an additional ingredient in it.

To elaborate a little. If A-A were simply one more, albeit very general rule of rational conduct, then a subject would be readily conceivable who had certain perfectly intelligible beliefs and desires but, because she had no intention to conform to A-A, routinely failed to act on them, performing instead in ways that her lower level beliefs and desires made no sense of. But that is not readily conceivable. A tendency to action upon beliefs and desires in a manner they rationally explain enters primitively into the conditions for *having* the appropriate beliefs and desires in the first place. A subject for whom akratic performance was the norm would so conduct herself as to undermine any basis for the ascription to her of any determinate set of beliefs and desires. What a subject does who is following a rule, personally or sub-personally, is subject to a conceptually unlimited degree of performance error, (even if it will eventually become a very implausible explanation of what is going on to continue to credit her with the continuing attempt to follow that rule and to look for the causes elsewhere.) But A-A does not allow for an unlimited degree of performance error: there is no such thing as having appropriate groupings of beliefs and desires yet consistently, though free to do so, failing to act in the ways they would rationalise. The very hypothesis of such a pattern takes the subject beyond the pale of rational agency and cuts the ground from under her possession of beliefs and desires in the first place. A-A is a norm of rational action not in the sense of something at compliance with which rational subjects somehow aim—or sub-personally aim, whatever that might mean—but in the sense that it is only by activity that allows of description as, broadly, in compliance with it that a subject accomplishes intentionality of any kind, only thereby that she gets passage into the space where it is *possible* to have aims and act on them.

So, what of it? Well, I have argued that inference cannot in general be constrained by the Taking Condition. But of course the Taking Condition is sometimes satisfied: sometimes a thinker moves to a conclusion because he does indeed take it that it is supported by his other acceptances. Let us set those cases

aside and call inference that is exceptional to the Taking Condition *basic inference*.⁸ I am suggesting that we compare basic inference—the kind of mental transition possibly illustrated by Boghossian’s Hide-and Seek example, but which, I have suggested the regress teaches us, must also be an ingredient at some level even in fully explicit personal inferential processing—with non-akratic action: let us think of basic inference as a form of, so to say, *basic mental action*. I have noted that if the comparison is acceptable, then there is no more need for a subject to have beliefs, or other forms of intentional registration, about a basic inference in order for her moving to its conclusion to *count* as inference than a rational agent needs to have beliefs about A-A, or more specifically about the sufficiency of her beliefs and desires to rationalise her action, in order for that action to count as action on those beliefs and desires, and to that extent rational. In brief: just as movement in accordance with A-A is constitutive of rational action, so movement in accordance with basic rules of inference is constitutive of rational thought.

This suggests—this is the third point I advertised—that we look at the role of *rules of inference* in basic inference on the model of the role of A-A in rational action. Thus it is not that some kind of sub-personal acceptance of, say, *modus tollens* underwrites and rationalises—even though sub-personal—the intelligent movements of thought of Boghossian’s Hide-and-Seeking children. Rather *modus tollens* stands to Luke’s performance exactly as A-A stands to it. He wants to find his friend and believes he is not behind that tree. That’s a complete rational explanation of why he rushes off to look elsewhere, without the need for any form of acceptance, personal or sub-personal, of A-A. Likewise, he sees the bicycle leaning against the tree and believes his friend would not hide in a spot where he had left his bicycle. That’s a complete rational explanation of why he believes his friend is elsewhere, without needing to advert to any form of acceptance, personal or sub-personal, of *modus tollens*.

I am under no illusion that these remarks do anything more than point in a direction. There is a great deal more to think about, but here are six points by way of a take-home:

- (i) Inference is, *au fond*, basic mental action: the formation of acceptances for reasons consisting of other acceptances.
- (ii) Such action no more requires control by states that register the sufficiency of the relevant reasons than action in general requires the presence, in the practical syllogism that rationalises it, of a state of the agent registering the sufficiency of the reasons supplied by his relevant beliefs and desires.
- (iii) The role of rules of inference in basic inference is not that of principles which are followed—personally or sub-personally, whatever might be the best account of that,—but that of constitutive norms, comparable to the role of A-A in rational action. They are constitutive in the sense that conformity to

⁸ Boghossian has often called this “blind inference”, and treated it as a special case of a more general phenomenon of blind rule-following. But I dislike this terminology, for the reasons given earlier: rule-following has better be subject to the Taking Condition, and ‘blind’ following can therefore connote only the following of a rule that is merely sub-personally registered. Whereas what I am now calling basic inference is not that.

them sets a defeasible but prevailing standard for the very possession of the acceptances among which they mandate movement. It is accordingly unhappy, independently of the Regress, to view them as rules which the rational (sub-personally) follow or tacitly know. Their contribution to the rationality of performance that accords with them is not mediated by states, at whatever level, of the subject that carry them as contents.

- (iv) A crucial corollary of point (iii) is that it is incorrect to view the patterns of acceptance-formation and revision that such norms require as the manifestation of (mere) *dispositions* to rational acceptance. To anything pointfully described as a disposition corresponds a possible pattern of activity that constitutes the lack of it: a pattern in which the distinctive manifestation of the disposition is at best irregularly triggered under the relevant conditions. But there is no such pattern associated with constitutive norms. The propriety of describing the states or actions among which they prescribe a rational connection in the very terms required to underwrite such a pattern of irregularity is undermined by the very hypothesis of the irregularity.
- (v) Any account along these lines needs to be supplemented by an accommodation of the idea that basic inferences may assume patterns that are *unsound*; human children, for example, and some adults, incline to Deny the Antecedent. This is a difficulty for the Simple proposal. When an inference is basic, so that the Taking Condition does not operate, but is an instance of a flawed but repeated pattern, a readiness to make it will hardly be a constraint on a thinker’s acceptances in the manner of the analogy with A-A. So what in such a case makes it an inference—constitutes the acceptance of its premises as the actual reasons that the thinker has for the move? I have said nothing about this, and mention it here only to show that I am aware of it.
- (vi) But we *would* know what to say about that if we knew what to say about what it is to act for certain specific reasons. The Simple proposal is completely invested in the good-standing of that notion. But it is a notion that, although it wears causality on its sleeve, has proved resistant to Naturalistic account. Thus the overall orientation in our thinking about inference that I have been promoting in the second half of these remarks therefore promises little to qualify Boghossian’s concluding pessimism about the prospects for our achieving a naturalistic understanding of inference.

Reference

- Coliva, A. (Ed.). (2012). *Meaning, knowledge and mind Essays in honour of Crispin Wright*. Oxford: Oxford University Press.