JOSÉ L. ZALABARDO

BONJOUR, EXTERNALISM AND THE REGRESS PROBLEM*

ABSTRACT. In this paper I assess the two central ingredients of Laurence BonJour's position on empirical knowledge that have survived the transition from his earlier coherentist views to his current endorsement of the doctrine of the given: his construal of the problem of the epistemic regress and his rejection of an internalist solution to the problem. The bulk of the paper is devoted to a critical assessment of BonJour's arguments against externalism. I argue that they fail to put real pressure on externalism, as they rely on a highly questionable conception of epistemic rationality and responsibility. Then, more briefly, I take issue with BonJour's endorsement of the *irrelevance thesis*—the claim that even if externalism were true it would not offer a satisfactory solution to the epistemic regress problem. I contend that he is not entitled to subscribe this thesis unless he is prepared to abandon his construal of the problem.

1. INTRODUCTION

Laurence BonJour's views on the structure of empirical knowledge have recently undergone a notorious transformation. His writings in the late 1970s and early 1980s provided a sophisticated defence of the coherence theory and an incisive attack on the version of foundationalism based on the doctrine of the empirically given.[1] However, in his recent work he has abandoned the former in favour of the latter.[2] This change of mind hasn't been brought about by a reassessment of the intrinsic merits of the doctrine of the given. Rather, BonJour has abandoned coherentism because he has come to the conclusion that some of the objections that it faces are fatal.[3] And this has lead him to embrace the doctrine of the given because he regards it as the only plausible alternative to coherentism.

The need to choose between coherentism and givenism is forced on BonJour by two important aspects of his earlier position that he still subscribes. The first is a construal of the problem of the epistemic regress according to which avoiding scepticism requires choosing between the coherence theory, the doctrine of the given and an externalist account of epistemic justification. The second is his conviction that the externalist solution to the problem doesn't work.

Synthese (2006) 148: 135–169
DOI 10.1007/s11229-004-2245-2

My goal here is to assess these two aspects of BonJour's position. The bulk of the paper is devoted to assessing his rejection of externalism. I shall contend that the arguments that he levels against externalism don't pose a serious threat to this view. Then I shall turn more briefly to considering the bearing of my discussion of externalism on the epistemic regress problem. I shall contend that, in order to uphold his conviction that externalism doesn't offer a satisfactory solution to this problem, BonJour would have to abandon his cconstrual of it.

2. THE EPISTEMIC REGRESS PROBLEM

Broadly speaking, the epistemic regress problem can be traced back to the ancient sceptics. It is the problem posed by the fact that our attempts to adduce evidence for our beliefs are bound to generate an infinite regress or a circle, or to end with an assumption for which no evidence is provided. These three possibilities correspond to three of the Five Modes that Sextus Empiricus attributes to "the later sceptics" and Diogenes Laertius, in a parallel account, to "Agrippa and his school".[4] Nevertheless, in spite of its long history, the nature of the problem is by no means clear, and even the claim that there is so much as a problem is open to question. Even if we concede that our attempts to defend our beliefs from the challenge of a persistent sceptic could only have one of these three outcomes, we still need an explanation of how, if at all, our epistemic situation is compromised by this circumstance.

BonJour endorses a construal of the epistemic regress problem from which these questions receive precise answers. I am going to characterise this construal of the problem as an argument for the negative epistemic conclusion that the trilemma is supposed to enjoin. The conclusion of the argument is that we can have no knowledge, because none of our beliefs can have the feature that confers on a true belief the status of knowledge (even if the difficulties posed by Gettier cases do not arise).[5] I shall use the term *warrant* to refer to this feature, making no assumptions concerning its nature.[6]

On the construal of the argument that BonJour endorses, this conclusion is reached with the help of two assumptions, expressing a twofold connection between warrant and evidence. The first is the claim that in order for a subject's belief in a proposition to have warrant, the subject has to have adequate evidence for the proposi-

tion. I shall refer to this principle as the *evidential constraint on warrant* (EC). The second is an account of what adequate evidence consists in, according to which a subject has adequate evidence for a proposition she believes only if she has other warranted beliefs whose truth she regards as making the truth of the target proposition likely.[7] I shall refer to this principle as the *inferential account of adequate evidence* (IA).

With EC and IA in place, the trilemma can be brought to bear on the possibility of knowledge. If p is a proposition that a subject S believes at a time t, we can easily define the notion of an *evidential tree* for p, for S at t, using the relation, defined on the set of propositions that S believes at t, that a proposition ϕ bears to a set of propositions Γ when S believes at t that the truth of the elements of Γ would make the truth of ϕ likely. The problem raised by EC and IA is that they entail that every branch of every evidential tree will have one of the following three shapes. It will be infinite, it will be circular, i.e., it will end with a proposition that has already appeared further down the branch, or it will end with a proposition that S believes without warrant.

If we define a *standard evidential tree* as one with no branches of any of these three shapes, we can say that EC and IA entail that there are no standard evidential trees. This result doesn't quite yield the conclusion of the argument. It follows from EC and IA that warrant requires an evidential tree in which no branch ends with an unwarranted belief, and one could argue that it is not humanly possible to believe infinitely many propositions, and hence that an evidential tree can't have infinite (non-circular) branches.[8] But it is certainly possible for an evidential tree to contain circular branches, and EC and IA by themselves do not rule out the possibility of warrant resulting from such a tree. Hence, to yield the intended conclusion, the argument needs an additional assumption to the effect that this is not a possibility. I shall refer to this assumption as the *non-circularity principle* (NC).

For a subject with finitely many beliefs, EC, IA and NC entail that none of her beliefs has warrant. I want to suggest that this argument provides a fairly accurate characterisation of the construal of the epistemic regress problem endorsed by BonJour, as well as by other contemporary epistemologists.[9] I shall refer to it as the *contemporary regress argument*.

Since, subject to the finitude proviso, the impossibility of warrant is deductively entailed by EC, IA and NC, vindicating the possibility

of warrant requires giving up one of these three premises. These three strategies correspond to the three positions in epistemology from which, according to BonJour, we are forced to choose. In the 1980s, BonJour held both EC and IA, and sought to vindicate the possibility of warrant in the face of the regress argument with his coherentist account, which licensed the rejection of NC. According to coherentism, a subject's belief in a proposition can have warrant even if all of the subject's evidential trees for the proposition have circular branches. Now that the objections raised against coherentism have convinced him to abandon the position, he has to accept NC. And since he still holds EC, his only chance of vindicating warrant is to reject the universal applicability of the inferential account of adequate evidence, arguing that in some cases we can have adequate evidence for a proposition we believe which doesn't take the form of other warranted beliefs. In these cases there is "an internally available reason why [the belief] is likely to be true without that reason depending on any further belief or other cognitive state that is itself in need of justification".[10] As BonJour is well aware, his new position faces very serious obstacles, but after abandoning coherentism he is left with no alternative, given that he remains as convinced as in his coherentist days that the evidential constraint is non-negotiable.

3. THE EXTERNALIST SOLUTION

But support for the evidential constraint is by no means unanimous. It is a fairly direct consequence of externalist accounts of knowledge that in some circumstances a subject's belief that *p* can have warrant even if she has no adequate evidence for *p*. These views were initially developed as attempts to solve problems with which we are not concerned here, but in 1973 David Armstrong explicitly put forward an account of knowledge along these lines as an attempt to defend the possibility of knowledge from the challenge of the contemporary regress argument.

Armstrong's solution arises from an account of knowledge for a class of beliefs which he calls 'simple judgments of perception'. On this account, if a subject *a* believes that *p* (Bap), and this belief is a true judgment of perception, then it will have the status of knowledge if there is "*a law-like connection* between the state of affairs Bap and the state of affairs that makes '*p*' true such that, given Bap, it must be the case that *p*".[11]

Thus on Armstrong's account of knowledge for simple judgments of perception, the existence of this law-like connection is a sufficient condition for warrant. But the existence of this connection wouldn't provide the subject with adequate evidence for her belief, since she might not even believe that the connection exists. It follows that, on Armstrong's account, the evidential constraint on warrant doesn't apply to simple judgments of perception, since they could be warranted even if the subject didn't have adequate evidence for them. We can now invoke this outcome to save the possibility of warrant from the threat of the regress argument. Warrant is made possible for any beliefs with evidential trees whose branches end with simple judgments of perception that sustain the relevant law-like connections.

This strategy for resisting the conclusion of the contemporary regress argument can be underwritten by accounts of knowledge that differ substantially from Armstrong's proposal. All that's required is a sufficient condition for warrant that can in principle be satisfied by a belief even if the subject doesn't have adequate evidence for it. In what follows, I'd like to focus my discussion, not on Armstrong's proposal, but on the account of warrant known as *process reliabilism*, to which BonJour himself devotes most of his attention.[12] The basic idea of process reliabilism is that whether a belief has warrant depends on the reliability of the cognitive devices which have played a role in its formation (or in sustaining it). To have some plausibility, an account of warrant along these lines would have to exhibit a considerable degree of complexity. However, for our immediate purposes, all that matters is that the account entails that if a belief has been formed with a reliable procedure, then it has warrant, at least in some circumstances in which the subject might not have adequate evidence for the belief. By *(process) reliabilism*, I shall refer to any account of warrant which entails a principle along these lines.

The reliability of the procedure with which a belief has been formed wouldn't by itself furnish the subject with adequate evidence for the belief. A warranted belief in the reliability of the procedure would have this effect, but a belief could have been reliably formed even if the subject doesn't believe that this is so, let alone believe it with warrant. Hence process reliabilism, like Armstrong's proposal, would license violations of the evidential constraint, thus blocking the regress argument. Warrant would be made possible for those beliefs that obtain it from the reliability of their formation process, and these warranted beliefs could then end the branches of evidential trees which confer warrant on further beliefs.

BonJour's assumption that his current position is the only alternative to his previous one presupposes that the externalist strategy for dealing with the contemporary regress argument is not available. Needless to say, BonJour doesn't merely assume this. In Chapter 4 of *The Structure of Empirical Knowledge*, he provided one of the most forceful attacks on externalism to date.[13] My goal in the remainder is to clarify and assess BonJour's anti-externalist arguments. My conclusions will be largely negative. I shall argue first that his attack on externalism is far from irresistible. Then I shall contend, more briefly, that his attitude to the externalist solution to the regress problem is incompatible with the construal of the problem that he endorses.

4. INTUITIONS

BonJour's attack on reliabilism is based on a series of well-known thought experiments involving subjects with the power of clairvoyance. The central case in BonJour's argument is that of Norman:

> Norman, under certain conditions which usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, although he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power under circumstances in which it is completely reliable. (p. 41)

BonJour clearly believes that our intuitions regarding this case undermine reliabilism. What is much less clear is the precise route by which he expects this conclusion to be reached.

A particularly straightforward strategy for adjudicating the contest between the evidential constraint and reliabilism is by appeal to our intuitions concerning the ascription of knowledge in specific circumstances. For this purpose, we would need to concentrate on cases in which a belief satisfies the sufficient condition for warrant postulated by reliabilism, but not the necessary condition for warrant laid down by the evidential constraint, i.e., cases of reliably formed belief for which the subject has no adequate evidence. If intuition regarding these cases dictated that the beliefs in question have warrant, this could be taken as evidence against the evidential constraint and, less directly, for reliabilism. If, on the contrary, intuition dictated that these beliefs lack warrant, we could take this result as

undermining reliabilism and, less directly, as supporting the evidential constraint. I shall refer to this strategy as the *intuitive approach*. One way to construe BonJour's argument against externalism is as an application of the intuitive approach. My goal in this section is to articulate this construal and to assess the resulting anti-externalist argument.

One difficulty faced by the intuitive approach is that it is remarkably hard to find real cases which both parties agree to describe as satisfying the reliable-formation condition but not the evidential constraint. The reliable-formation condition would be satisfied by cases in which a belief was formed as a result of the operation of a basic cognitive device that is reliable, such as properly functioning perceptual devices, or introspection. But the problem with these cases is that it is open to the proponent of the evidential constraint to argue that the subject has adequate evidence for these beliefs, typically taking the form of a belief identifying the procedure with which the target belief has been formed and a belief in the reliability of this procedure. Take, e.g., Alicia's belief that a train is approaching the platform, formed by her highly reliable visual device for the detection of events of this kind. If the reliabilist tried to appeal to our intuition that Alicia knows that a train is approaching the platform as supporting her position, the advocate of the evidential constraint could reply that it would be illegitimate to interpret the intuition in this way. She would explain the intuition instead as resting on a background assumption to the effect that Alicia has adequate evidence for her belief. For in normal conditions we would expect Alicia to believe that her belief that a train is approaching the platform was formed with the relevant visual device, and that this device is reliable. Furthermore, we would expect these beliefs of hers to be warranted. Clearly, if this assumption is playing a role in generating our intuition, it cannot be counted as lending support to reliabilism or undermining the evidential constraint.

We could of course stipulate that Alicia doesn't have adequate evidence for her belief, but it is questionable whether this stipulation would be effective in insulating our intuitions from the background assumption that the stipulation is meant to suspend. It is open to the proponent of the evidential constraint to argue that if we still have the intuition that Alicia knows that a train is approaching the platform after stipulating that Alicia doesn't have adequate evidence for her belief, the reason is that the stipulation hasn't succeeded in neutralising the effect on our intuitions of the background assumption that she does

have adequate evidence. Let me refer to the possibility that our stipulations concerning a thought experiment might not have the intended effect on our intuitions as the problem of *intuition contamination*.

In sum, the attempt to use the intuitive approach to adjudicate the contest between reliabilism and the evidential constraint can be expected to generate the following dialectic. First, the reliabilist will claim that the method supports her position, on the grounds that we have stable intuitions to the effect that when a belief has been formed by one of our basic cognitive devices, and the device is reliable, the belief has warrant. But even if she concedes that we have these intuitions, the evidentialist can object to the reliabilist's attempt to use them in her favour, arguing that they are to be explained instead by reference to a background assumption to the effect that in such cases the subject typically has adequate evidence for her belief. The reliabilist might then argue that these intuitions survive the stipulation that the subject doesn't have adequate evidence, but the evidentialist could question the significance of this fact, explaining it as a case of intuition contamination.

The only way forward for the intuitive approach appears to be to find cases in which we can effectively stipulate that the subject lacks adequate evidence for her belief. If intuition favours the ascription of warrant in these cases, the evidentialist won't be able to appeal to intuition contamination to rescue her position. If, on the contrary, intuition dictates that in these cases we shouldn't ascribe warrant, the evidentialist could claim that her objections to the intuitions that appear to support reliabilism have been vindicated.

BonJour's introduction of his clairvoyance cases appears to be motivated by these reflections. His goal is to avoid the unfair advantage that intuition contamination would give to the reliabilist in more familiar cases. "Cases involving sense-perception and introspection", he tells us, "are […] not very suitable for an intuitive assessment of externalism, since one central issue between externalism and other […] views is precisely whether in such cases a further basis for justification beyond the externalist one is typically present" (pp. 37–38).[14] To avoid this difficulty, we need to concentrate on cases "for which it will be easier to stipulate in a way which will be effective on an intuitive level that *only* the externalist sort of justification is present" (p. 38).

Norman's case might seem to fit the bill perfectly. His belief that the President is in New York City exhibits the features that are required for the application of the intuitive approach. On the one

hand, the belief has been formed by a reliable procedure – Norman's powers – thereby satisfying the sufficient condition for warrant postulated by the reliabilist.[15] On the other, the evidential constraint is not satisfied, as Norman doesn't have adequate evidence for his belief. BonJour considers two versions of the case.[16] In the first version, Norman takes himself to have evidence for his belief that the President is in New York City, since he believes it to have been formed with his clairvoyant faculty and he believes that this faculty is reliable. However this evidence is not adequate, as Norman's belief in the reliability of his clairvoyant faculty lacks warrant. In the second version, Norman doesn't even take himself to have evidence for his belief, since he doesn't believe that he has the power of clairvoyance.

Furthermore, applying the intuitive approach to this case won't raise the difficulties that we encounter when we consider more familiar cognitive devices such as sense perception. Reliabilists won't obtain an unfair advantage from intuition contamination in this case, since we have no irresistible inclination to assume that Norman has adequate evidence for his belief. Hence if intuition dictated that Norman doesn't know that the President is in New York City, the evidentialist could claim that, once it is properly applied, the intuitive approach adjudicates in favour of her position. I am not going to discuss whether intuition actually dictates that Norman's belief lacks warrant. I am going to assume that it does,[17] and argue that it is open to the reliabilist to resist the thought that this outcome puts real pressure on her position.

My proposal is that the reliabilist could contend that in moving from sense perception and introspection to clairvoyance, we may have removed one potential source of contamination – but only to replace it with another one which distorts our intuitions in the opposite direction. The problem is that many of us have a deep-seated conviction that clairvoyance doesn't exist, that its possibility is ruled out by some central elements of our world-view, and that evidence for this assessment is readily available. Of course, in the Norman case we stipulate that these assumptions are false, but it is open to the reliabilist to argue that these stipulations might not succeed in insulating our intuitions from the assumptions that they are meant to neutralize.

There are two particular aspects of the case where contamination might be suspected. The first one is the claim that Norman has real clairvoyant powers. As some authors have suggested, many of us find

it very hard to imagine a world in which clairvoyance is possible. If this is so, there is reason to suspect that our negative verdict on Norman's belief is informed by our inability to suspend disbelief in the existence of clairvoyance. According to this diagnosis, the reason why we don't think Norman knows is not that we don't take reliable formation to be sufficient for warrant, but that we can't bring ourselves to assume that his belief has been formed with a reliable mechanism.[18]

The second feature of the case that might raise the suspicion of contamination is the claim that Norman doesn't have adequate evidence for the proposition that he doesn't have clairvoyant power. Notice that if he had such evidence, then reliabilism could easily accommodate the intuition that Norman doesn't know that the President is in New York City. On Goldman's version of the position, the definition of warrant incorporates a *non-undermining provision*, according to which a reliably formed belief wouldn't have warrant if the subject had evidence against the reliability of the procedure with which it has been formed.[19] Hence, if Norman had evidence against the proposition that he has clairvoyant powers, the intuition that his belief doesn't have warrant would be in line with Goldman's account. BonJour accepts this, and the stipulation that Norman doesn't have any evidence or reasons against the existence of clairvoyance or against the hypothesis that he possesses it is explicitly designed to block this reliabilist rejoinder. Goldman's original response to the case was to maintain that Norman does have evidence against the hypothesis that he has clairvoyant powers, and hence that his non-undermining provision can deal with the case,[20] although he has later seemed less confident in this reply.[21] However, even if we concede to the evidentialist the right to stipulate that Norman has no evidence against clairvoyance, we can still question the power of this stipulation to protect our intuitions from contamination, given that we are inclined to assume that a standard belief system provides adequate evidence for the proposition that clairvoyance doesn't exist.[22]

In sum, the reliabilist can resist the anti-reliabilist application of the intuitive approach based on Norman's case with the same tools that the evidentialist deployed to resist the pro-reliabilist application. The anti-reliabilist significance of our intuitions in Norman's case, like the pro-reliabilist significance of our intuitions in the case of sense perception, is tainted by the suspicion of contamination.[23]

5. BONJOUR AND THE INTUITIVE APPROACH

In Section 4 I have developed a construal of BonJour's anti-externalist case as an application of the intuitive approach. I have argued there that using Norman's clairvoyance thought experiments in this way wouldn't result in a very powerful anti-reliabilist argument. In this section I am going to question the attribution of this argument to BonJour.

BonJour manifests his attitude to an anti-reliabilist application of the intuitive approach when he is evaluating the attempt to deploy this strategy in defence of reliabilism. He characterises the starting point of this argument in the following terms:

> The basic factual premise of the argument is that in many cases which are commonsensically instances of justified belief and of knowledge, there seem to be no justifying factors present beyond those appealed to by the externalist. (p. 52)

A proponent of an anti-reliabilist application of the intuitive approach can be expected to reject this premise – to contend that in cases (like Norman's) in which no justifying factors are present beyond those appealed to by the reliabilist, intuition dictates that the beliefs in question don't have the status of knowledge. However, BonJour's assessment of this premise is completely different:

> Though the precise extent to which it holds could be disputed, in the main the initial factual premise of this argument must simply be conceded. Any nonexternalist account of empirical knowledge that has any plausibility will impose standards for justification that many commonsensical cases of knowledge will fail to meet in any full and explicit way. (pp. 52–53)

BonJour then goes on to question the strength of the support that reliabilism would receive from this premise, arguing that an account of warrant might in principle be correct even if it forces us to revise some of our intuitive ascriptions of knowledge. A defence of reliabilism that takes these intuitions as the final arbiter for the correctness of an account of warrant would rely on the following additional premise:

> [...] that the judgments of common sense as to which of our beliefs qualify as knowledge are sacrosanct, that any serious departure from them is enough to demonstrate that a theory of knowledge is inadequate. (p. 53)

But this premise, BonJour tells us,

> [...] seems entirely too strong. There seems in fact to be no basis for more than a quite defeasible *presumption* (if indeed even that) in favor of the correctness of common sense. (Ibid.)

Although BonJour's intention is to undermine a defence of reliabilism along the lines of the intuitive approach, it is clear that his position would undermine to the same extent an attempt to use this strategy to attack reliabilism.

In light of these considerations, we have to conclude that BonJour takes a very dim view of the attempt to defend the evidential constraint by applying the intuitive approach to his clairvoyance cases. He thinks that conflict with our intuitive warrant ascriptions wouldn't seriously undermine an account of warrant, and even if it did, the proponent of the evidential constraint wouldn't be able to use this circumstance to her advantage, as reliabilism provides a better match for our intuitive warrant ascriptions than the evidential constraint. The argument that uses the intuitive approach to undermine reliabilism can't be BonJour's.

6. EPISTEMIC RATIONALITY AND RESPONSIBILITY

This raises the question of how we should construe the reasoning in which BonJour appeals to the clairvoyance cases to undermine reliabilism. He tells us that his goal in presenting these cases is to support and develop the following intuitive difficulty with externalism:

> according to the externalist view, a person may be highly irrational and irresponsible in accepting a belief, when judged in light of his own subjective conception of the situation, and may still turn out to be epistemically justified [...] (p. 38)

This aspect of the argument is emphasized in a later passage, where he describes his anti-externalist argument as putting forward the following conception of epistemic rationality: ''of such rationality as essentially dependent on the believer's own subjective conception of his epistemic situation'' (pp. 49–50).

These passages suggest an anti-reliabilist case based on the contention that reliabilism ascribes warrant to a belief in cases in which the subject is epistemically irrational and irresponsible in holding it (relative to his own subjective conception of the situation). According to this line of thought, reliabilism would have to be rejected on the grounds that this consequence is unacceptable. An argument against reliabilism along these lines would invoke the following two premises:

1. Reliabilism entails that a belief can have warrant even if the subject is epistemically irrational and irresponsible in holding it (relative to his own subjective conception of the situation).

2. A belief can't have warrant if the subject is epistemically irrational and irresponsible in holding it (relative to his own subjective conception of the situation).

I want to suggest that this is the main argument that BonJour is advancing against reliabilism, and that the role of the clairvoyance thought experiments in his discussion is to provide support for these premises. I shall consider each premise in turn.

Before we proceed, let me mention briefly a prominent line of reasoning that I shall not be considering here. One could object to this argument on the grounds that the notions of responsibility and rationality can't be applied to belief, because these notions can only be applied to items that are under the voluntary control of the subject, and belief doesn't satisfy this condition: we cannot decide what to believe. I shall not try to assess this line of reasoning here. I only want to point out that its success is open to question, and hence that the existence of other weaknesses in BonJour's reasoning would not be of merely academic interest.[24] In what follows I am going to waive this difficulty and assume for the sake of the argument that it makes perfect sense to ascribe to belief rationality and responsibility, as well as their opposites.

7. PREMISE 1

The core of BonJour's case for Premise 1 is the thought that the principle is instantiated by Norman's case. On the one hand, the reliabilist is committed to ascribing warrant to Norman's belief that the President is in New York City. On the other hand, Norman is epistemically irrational and irresponsible in holding this belief (relative to his own subjective conception). As we saw in Section 4, the first of these claims has not gone unchallenged. Nevertheless it is the second claim that I want to consider here. My goal is to assess the claim that, in holding his belief about the President's whereabouts, Norman is epistemically irrational and irresponsible (relative to his own subjective conception).

For this purpose we need to clarify the notion of rationality and responsibility that figures in the claim. The first point we need to note is that according to BonJour, what distinguishes epistemic rationality or responsibility from other forms of rationality or responsibility is its relation to the cognitive goal of truth. A precise construal of this goal raises familiar problems, but for our purposes we just need to assume

that it is a goal that would be promoted, for any given proposition, by believing it if it is true, and not believing it if it is false.

I propose then to think of epistemic rationality and responsibility as instrumental, or means-ends rationality, relative to the goal of believing a proposition if and only if it is true. On this construal of the notion, the epistemic rationality of holding a belief will depend on the extent to which doing so promotes the cognitive goal of truth.

Now, the instrumental rationality of an action can be assessed from several points of view. One possibility is to assess it from an objective perspective, in terms of the objective likelihood of the action promoting the relevant outcome. Alternatively, we can assess it from the point of view of a certain conception of the relevant situation, independently of whether this conception is correct. The assessments of instrumental value that we would obtain from different conceptions of the situation can be expected to differ from each other, and the assessments that would result from an incorrect conception can be expected to differ from those that we would obtain from the objective point of view.

Thus, on this construal, BonJour is claiming that from the point of view of Norman's own conception of his situation, holding his belief that the President is in New York City is of little instrumental value, relative to the cognitive goal of truth. But the instrumental value of holding a belief relative to the cognitive goal of truth will depend on how likely it is that the belief is true. Hence, on my construal, BonJour's claim is that from the point of view of Norman's own conception of his situation, his belief that the President is in New York City is unlikely to be true.

In order to assess this claim, we need to reach a more precise understanding of the notion of "Norman's conception of his own epistemic situation". As a first approximation, we can say that how likely it is, from the point of view of a subject's conception of her situation, that an action will promote a goal, is determined by how likely she believes it to be that the action will promote the goal. This thought yields the following account of subjective rationality and responsibility:

R1 From the point of view of a subject's conception of her situation, it is rational and responsible for her to perform action A in pursuit of goal G just in case she believes that A is likely to promote G.[25]

In the particular case of epistemic rationality and responsibility, the account would go as follows:

ER1 From the point of view of a subject's conception of her epistemic situation, it is epistemically rational and responsible to believe that *p* just in case she believes that *p* is likely to be true.

Notice that it follows from this account that for a subject who believes that *p*, holding this belief would be epistemically irrational and irresponsible from the point of view of her conception of her situation only if she didn't believe that *p* is likely to be true. One is entitled to wonder whether this is a coherent possibility–whether there is a possible cognitive state that we would want to describe by saying that the subject believes that *p* but doesn't believe that *p* is likely to be true. But even if we accept this as a genuine possibility, it is hard to see how BonJour could exploit it to defend the epistemic irrationality and irresponsibility of Norman's belief, since nothing in his description of the case indicates that Norman fails to believe that it is likely that the President is in New York City. The point is particularly clear in the first of the two versions of the case that BonJour considers, on which I shall focus my discussion in the first instance. In this version, "Norman *believes* himself to have clairvoyant power even though he has no justification for such a belief", and this belief "contributes to his acceptance of the belief about the President's whereabouts" (pp. 41–42). Clearly, if Norman believes that he has clairvoyant powers, he will believe that the propositions that are sanctioned by this faculty, including the proposition that the President is in New York City, are very likely to be true. I think we can conclude that on this construal of the notion of the subject's conception of her situation, there are no grounds for accusing Norman of irrationality and irresponsibility.

Since BonJour cannot exploit to his advantage the possibility of believing that *p* while failing to believe that *p* is likely to be true, we can simplify matters by ignoring it. Thus in what follows I shall assume, for the sake of simplicity, that a subject believes that *p* if and only if she believes that *p* is likely to be true.

This would be the end of the story if the account of subjective rationality and responsibility expressed by R1 and ER1 were correct. But the account needs to be revised. The problem is that, as it stands, it allows a subject to achieve rationality and responsibility by simply refusing to face the facts. Suppose that Alex drinks one bottle of whisky a day in order to improve his health, and he actually believes that drinking one bottle of whisky a day is likely to have this effect,

but only because he has deliberately avoided exposure to medical opinion about the likely effects of this level of alcohol consumption. It seems to me that in these circumstances we might not want to say that from the point of view of Alex's conception of the situation, drinking one bottle of whisky a day is rational and responsible, relative to the goal of preserving his health, even though he believes that this policy is likely to promote the relevant goal. The same point can be made in the case of epistemic rationality. Suppose that Tracey believes that her husband liked the present she bought for him, but only because she hasn't asked him, although she believes that he would tell her the truth if she asked, but would otherwise act as if he had liked the present, whether he did or not. One could argue that from the point of view of Tracey's conception of the situation it is not rational for her to hold this belief, relative to the cognitive goal of truth.

These examples indicate that there are circumstances in which one can plausibly claim that the subject's beliefs do not determine what's rational and responsible for her to do in pursuit of a goal, from the point of view of her conception of her situation. I want to suggest that these are all cases in which the subject hasn't done her best, by her lights, to determine the likelihood that the course of action in question will bring about the relevant outcome. Alex may believe that one bottle of whisky a day will keep the doctor away, but he also believes that his doctor knows more about these things than he does, and he hasn't consulted her. Tracey may believe that her husband liked the present, but she also believes that she could easily check by asking him, and she has failed to do so. There are cognitive procedures that they haven't implemented, although they could have done so, which they think might have affected their assessment of the likelihood that the course of action in question will bring about the relevant goal.

We can accommodate this point by restricting the connection between the subject's conception of the situation and the beliefs she holds about it to cases in which she has done her best by her lights to determine the likelihood of the action in question promoting the relevant goal. This move would leave us with the following weaker principle:

R2 If a subject has done her best by her lights to determine whether action A is likely to promote goal G, then from the point of view of her conception of her situation, it is rational and

> responsible for her to perform action A in pursuit of goal G just in case she believes that A is likely to promote G.

If we apply this general principle to the epistemic case we obtain the following:

> ER2 If a subject has done her best by her lights to determine the truth value of a proposition *p*, then from the point of view of her conception of her epistemic situation, it is epistemically rational and responsible to believe that *p* just in case she believes that *p* is likely to be true.

And by virtue of our simplifying assumption, this is equivalent to the following principle:

> ER2* If a subject has done her best by her lights to determine the truth value of a proposition *p*, then from the point of view of her conception of her epistemic situation, it is epistemically rational and responsible to believe that *p* just in case she believes that *p*.

I want to suggest that this revised principle provides a very accurate constraint on the notion of the epistemic rationality and responsibility of a belief from the subject's point of view. According to it, if a subject believes that *p*, then this belief could only be epistemically irrational and irresponsible from the point of view of the subject's conception of her situation if she has failed to do her best by her lights to determine the truth value of *p*.[26]

This gives us a very precise picture of how BonJour would have to proceed to substantiate his claim that from the point of view of Norman's conception of his epistemic situation his belief that the President is in New York City is epistemically irrational and irresponsible. He would have to argue that Norman hasn't done his best by his lights to determine whether the President is in New York City. The most direct route to this conclusion would be to contend that, since Norman has gathered no evidence regarding the President's whereabouts beyond the testimony of his clairvoyant powers, he cannot have done his best by his lights to determine whether he is in New York City, and hence his epistemic rationality and responsibility won't be guaranteed by EC2*. However, this argumentative strategy faces serious obstacles. The problem is that whether you have done your best to attain a certain goal will depend, not only on what you have done, but also on what you could have done. Doing your best

can only require doing things that you are in a position to do. In particular, doing your best to determine the truth value of a proposition can only require obtaining evidence of a certain standard if this is something you are in a position to achieve. Hence, depending on the circumstances, a subject might have done her best to determine the truth value of a proposition she believes even if her evidence for it is not superior to Norman's evidence for the proposition that the President is in New York City.

Let me illustrate this point with an example. Clara, after wandering for a couple of hours in the Marrakech souk, decides to return to her hotel. She can't reconstruct the route that has taken her to her current location, and she can't identify any visual clues, but she is firmly convinced that the hotel lies to the right of where she is standing. Now, the situation might be such that Clara is capable of gathering evidence that would support or undermine the verdict of her sense of direction. She might have a map or a GPS, or there might be locals in the vicinity that could give her directions. In this case, if she neglects to use these sources of evidence, she can be accused of failing to do her best to determine the truth value of the proposition that the hotel lies to her left. But Clara could also find herself in a situation in which none of these procedures is available to her. She might have no map or GPS, and there might be no one around who might give her directions, etc. In these circumstances, it would be wrong to accuse her of failing to do her best to determine the truth value of the proposition that the hotel lies to her left, even if her sense of direction is all she can adduce in its support. Furthermore, when the question is whether she has done her best *by her lights*, what matters is not whether she has failed to use sources of evidence that are as a matter of fact accessible and relevant, but whether she has failed to use sources of evidence that she believes to be accessible and relevant. In general, she could only be accused of failing to do her best by her lights to determine the truth value of the proposition if she has failed to consult a source of evidence of which she believes (a) that it is available to her (monolingual Clara might feel incapable of obtaining directions from the locals) and (b) that its verdict might make a difference to her views on the matter (distrustful Clara might trust the locals so little that their directions are unlikely to affect her conviction that the hotel lies to her left). The accusation would be ungrounded if there are no sources of evidence satisfying these conditions that she has failed to consult.[27]

We can now apply these points to Norman's case. We have seen that BonJour needs to argue that Norman hasn't done his best by his lights to determine the truth value of the proposition that the President is in New York City. To support this conclusion BonJour would need to contend that there is a source of evidence satisfying the conditions of accessibility and relevance that Norman has failed to apply to the proposition that the President is in New York City. But I can't see that there is any legitimate reason to conclude that this is the situation in which Norman finds himself. Nothing in BonJour's description of the case rules out the possibility that, of every potential evidence source that Norman is aware of but has failed to apply to the proposition in question, he believes either that he can't apply it (e.g., tapping the President's phone might make him change his mind, but he can't really do it) or that it wouldn't make a difference to his verdict (e.g., he could ask his son, but what does he know…). And if this is the situation, it follows from ER2* that Norman's belief is epistemically rational and responsible from the point of view of his conception of his situation.

A specific version of the line of reasoning we have just considered deserves special attention. There is a strong body of widely available scientific evidence against the existence of clairvoyance. Since Norman doesn't have evidence against the possibility of clairvoyance, he has failed to consider the verdict of science. But this means that Norman hasn't done his best to determine the truth value of his belief, and hence ER2* doesn't entail that his belief is epistemically rational and responsible.

I suspect that this line of reasoning is responsible for much of the intuitive appeal of BonJour's verdict. However, it should be clear that it rests on an illegitimate assumption. By failing to consider the scientific evidence, Norman would only have failed to do his best by his lights if he believed about this procedure that he could apply it in this case and that the result of so doing might make a difference to his ultimate verdict. But it is not part of the case that Norman has either of these beliefs. The details of the case could be filled in in such a way that Norman believes that he has no access to the relevant scientific evidence. And even if he believes that he could consider the scientific evidence if he wanted, he might think that his belief that he has clairvoyant powers is so strong (or that his belief in science is so weak) that whatever the scientists might say would not make a difference to his conviction. And if his belief is not the result of culpable self-deceit or negligence, in failing

to take scientific evidence into account he wouldn't be failing to do his best *by his lights*.

To be sure, someone who holds science in such low esteem might be open to criticism, and we might even want to accuse them of irrationality or irresponsibility. But if their disdain of science results from the honest and conscientious pursuit of the cognitive goal of truth, the irrationality or irresponsibility they might be accused of is not irrationality or irresponsibility from the point of view of their conception of the situation. They might be irrational and irresponsible from our point of view, or even from the objective point of view. But from their own subjective perspective they are not doing anything wrong. From their own point of view their rationality and responsibility would not be open to question.

In his discussion of the first version of Norman's case, BonJour suggests another line of reasoning which it is tempting to regard as the real source of his conviction that Norman's belief is epistemically irrational and irresponsible. The starting point of this train of thought is what BonJour sees as the obvious irrationality of Norman's belief that he has clairvoyant powers:

> But is it not obviously irrational, from an epistemic standpoint, for Norman to hold such a belief when he has no reasons at all for thinking that it is true or even for thinking that such a power is possible? (p. 42)

Thus BonJour declares Norman's belief in his clairvoyance irrational on the grounds that he has no evidence or reasons to support it. Notice that the irrationality of Norman's belief that the President is in New York City cannot be established in quite the same way. In this version of the case, Norman has evidence for his belief, since he takes it to be sanctioned by a highly reliable procedure. The problem in this case is that the evidence is not adequate. Norman has beliefs of which he believes that their truth would make it likely that the President is in New York City, but these beliefs lack warrant. As BonJour tells us, "this belief about his clairvoyance fails after all to have even an externalist justification" (ibid.). Thus this line of thought would derive the irrationality of Norman's belief that the President is in New York City from his lack of adequate evidence for it.[28]

I want to argue that this line of reasoning is unlikely to provide a significant contribution to BonJour's argumentative strategy. The problem is that it rests on the assumption that epistemic rationality requires adequate evidence. We can formulate the principle that is being invoked in the following terms:

Evidential constraint on epistemic rationality and responsibility: If a subject doesn't have adequate evidence for a proposition *p* that she believes, then from the point of view of her conception of her situation she is epistemically irrational and irresponsible in believing that *p*.

I don't want to discuss the intrinsic merits of this principle. I only want to remind the reader that BonJour's goal is to defend the evidential constraint on warrant against externalist accounts of the notion, which seek to solve the epistemic regress problem by rejecting the constraint. And it is hard to see what argumentative strength would accrue to an argument for the evidential constraint on warrant that uses as a premise the evidential constraint on epistemic rationality and responsibility. This line of reasoning would be effective just in case those who are inclined to reject the evidential constraint on warrant could be expected to be more reluctant to reject the evidential constraint on epistemic rationality and responsibility. But I see no reason to expect this. In fact, the rejection of the evidential constraint on epistemic rationality and responsibility can be supported independently of the fate of the evidential constraint on warrant – by the adoption of an account of epistemic rationality and responsibility along the lines of ER2.[29]

I have focused so far on the first version of Norman's case. However, my conclusions apply equally to the second version, in which Norman finds himself strongly inclined to believe that the President is in New York City, but has nothing to say in support of his belief, since, unlike in the first version, Norman doesn't believe that he has the power of clairvoyance. However, as we have seen, according to ER2*, a belief can in principle be rational and responsible from the point of view of the subject's conception of her situation even if she has no evidence for it. This will be so if the subject has done her best by her lights to determine the truth value of the proposition in question. And we cannot rule out the possibility that Norman, in the second version of the case, as well as in the first, has applied to the proposition that the President is in New York City every source of evidence of which he believes that he could apply it and that doing so might make a difference.

It seems to me that BonJour's case at this point trades on an equivocation when he suggests that "from [Norman's] standpoint, there is apparently no way in which he *could* know the President's whereabouts", that "there is no way, as far as he knows, for him to have obtained this information" and that Norman's belief amounts

to "believing things to which one has, to one's knowledge, no reliable means of epistemic access" (p. 42). On one reading of these passages, they portray Norman as believing that he has no way of acquiring information about the President's whereabouts. If this is what Norman believes, then it is hard to see how he could fail to believe also that the chances of his belief being true are no better than random, and, according to ER2, this would make his belief subjectively irrational. But nothing in BonJour's description of the situation indicates that Norman believes that he has no way of obtaining information about the President's whereabouts. All that follows from the description of the case is that Norman can't identify a source of information about this matter, and it is perfectly conceivable that Norman believes that he has a source of information about these matters even though he can't identify it. And in these circumstances Norman's belief that the President is in New York City might not be subjectively irrational.

So far in my discussion of epistemic rationality and responsibility I have argued that BonJour is not entitled to claim that from the point of view of Norman's own conception of his situation he is epistemically irrational and irresponsible in believing that the President is in New York City. But notice that it is open to BonJour to concede this point while still maintaining that Norman's case supports Premise 1. All that is required to establish Premise 1 is to show that it is possible that reliabilism ascribes warrant to a subjectively irrational belief, and although BonJour's description of the Norman case is compatible with a situation in which Norman would not be subjectively irrational, it is also compatible with a situation in which he is subjectively irrational. Suppose, for example, that he has failed to ask his wife, who works in the White House travel office and has always been willing to give Norman accurate details of the President's travel arrangements. Then the rationality of his belief would no longer be guaranteed by ER2*. But since reliabilism would confer warrant on Norman's belief independently of how these details are filled in, Norman's case still supports Premise 1.

This is a perfectly legitimate move, and it will be effective so long as the reliabilist wants to maintain that the subjective irrationality of Norman's belief would leave its warrant intact. However, the move would be ineffectual against a reliabilist who is prepared to give up this thought, arguing that the circumstances that would undermine the subjective rationality and responsibility of Norman's belief would also undermine its warrant. This line could be pursued from within

Goldman's version of reliabilism, since one could argue that the circumstances that would render Norman's belief subjectively irrational would also bring it under the non-undermining provision. Alternatively, the reliabilist could adopt a bolder strategy. She could simply modify her account of warrant by treating epistemic rationality and responsibility, construed along the lines of ER2, as an additional necessary condition for warrant. The resulting view would still be recognisably reliabilist, and, more importantly for our purposes, it would still sustain an externalist solution to the epistemic regress problem, by licensing violations of the evidential constraint on warrant. And, needless to say, this version of reliabilism would make Premise 1 immediately false, provided that, as I am suggesting, ER2 is correct.

8. PREMISE 2

Let's now turn briefly to Premise 2 of BonJour's argument against reliabilism. The first point to notice is that the claim stands in need of support. It is far from obvious that one cannot acquire knowledge "against one's better judgment". Thus, e.g., Alvin Plantinga has argued that fulfilment of one's epistemic duties is not necessary for warrant.[30] Plantinga illustrates his point by asking us to suppose that he is convinced that forming the belief that he is seeing something red will deprive him of any chance of epistemic excellence. He argues that this conviction leads him to acquire an epistemic duty not to believe that he is seeing something red. However, he has the standard doxastic inclinations and it is only by heroic effort that he manages to refrain himself from forming this belief in the presence of fire engines, etc. Then he describes the following situation:

> On a given morning I go for a walk in London; I am appeared to redly several times (postboxes, traffic signals, redcoats practising for a reenactment of the American Revolution); each time I successfully resist the belief that I see something red, but only at the cost of prodigious effort. I become exhausted, and resentful. Finally I am appeared to redly in a particularly insistent and out-and-out fashion by a large red London bus. "Epistemic duty be hanged", I mutter, and relax blissfully into the belief that I am now perceiving something red. […] in accepting [this belief] I would be going contrary to epistemic duty; yet could it not constitute knowledge nonetheless?[31]

The case can easily be adapted to support the rejection of Premise 2.

In any case, I am not going to engage here in a general discussion of Premise 2. I am going to concentrate instead on the argument that BonJour adduces it its support. He presents the argument in his dis-

cussion of a possible reliabilist rejoinder. Before presenting Norman's case (case 4) BonJour presents three other cases of reliable clairvoyants. What distinguishes these cases from Norman's is that these subjects have reasons for thinking either that the President is not in New York City (case 1), or that they don't have clairvoyant powers (case 2), or that clairvoyance is impossible (case 3). The reliabilist strategy that BonJour is discussing consists in accepting that the subjects in the first three cases lack warrant, while refusing to extend this verdict to Norman. Here is BonJour's objection to this position:

> But how can case 4 be successfully prised apart from the earlier ones? What the externalist needs at this point is a different account of *why* the beliefs in cases 1–3 are not justified, an account which does not invoke the notion of subjective irrationality, and hence does not extend readily to case 4. (p. 46)

In this passage, BonJour is taking for granted the subjective irrationality of Norman's belief, and arguing against the view that it might nevertheless have warrant. In other words, he is arguing for the claim expressed in Premise 2 that subjective irrationality is incompatible with warrant. BonJour argues for this conclusion on the grounds that subjective irrationality is the best explanation of why the subjects in cases 1–3 lack knowledge, and this explanation would not be available if subjective irrationality were compatible with warrant. BonJour devotes most of his efforts to arguing against Goldman's alternative explanation of the absence of warrant in cases 1–3, apparently thinking that the optimality of his preferred explanation could only be undermined by the existence of a better candidate. Here I want to pose a more basic objection to BonJour's claim. The hypothesis that the subjects in cases 1–3 are subjectively irrational wouldn't even be in the running for the post of the best explanation of their lack of warrant unless it actually explained this phenomenon. And it wouldn't explain this phenomenon unless it was true. Hence BonJour's argument for Premise 2 would collapse if we could reject the claim that the subjects in cases 1–3 are subjectively irrational. This is precisely what I propose to do. I am going to argue that nothing in BonJour's description of his first three cases indicates that their subjects are subjectively irrational.

My reasons for making this claim are strictly parallel to the reasons I adduced for the Norman case. Let's concentrate on case 1. What makes Samantha's belief subjectively irrational, on BonJour's account, is the fact that she maintains it

appealing to her alleged clairvoyant power, even though she is aware of apparently cogent evidence, consisting of news reports, press releases, allegedly live television pictures, and so on, indicating that the president is at that time in Washington, D C. (p. 38)

Let's consider what can be said about the epistemic rationality and responsibility of Samantha's belief from the point of view of ER2*. The most direct route to the conclusion that according to ER2* Samantha is epistemically irrational and irresponsible would involve arguing that she hasn't done her best by her lights to determine whether the President is in New York City. Then the epistemic rationality and responsibility of her belief would no longer be protected by ER2*. But on the basis of what we know about Samantha, this would require arguing that she has failed to consider the bearing of the "media" evidence on her belief that the President is in New York City, and that she believes that implementing this procedure might make a difference to her final verdict on the truth value of this proposition. It is not clear whether this second point is part of BonJour's description of the case – that would depend on whether the evidence appears cogent *to Samantha*. If, having done her best by her lights to determine the reliability of media evidence, she believes that it is not to be trusted, then by failing to take it into account she wouldn't be failing to do her best by her lights to determine whether the President is in New York City – even if the media evidence appears cogent to us, or if it appears cogent in some objective sense.

But even if this point is conceded it doesn't follow that Samantha hasn't done her best by her lights, since we are not entitled to assume that she hasn't considered the bearing of the media evidence on her belief that the President is in New York City. The only reason to assume this would be the thought that since the evidence appears cogent to her, if she had considered it she would have concluded that the President is unlikely to be in New York City, and hence, by our simplifying assumption, she wouldn't have the belief which by hypothesis she does have. But this reasoning rests on a mistake. From the assumption that the media evidence appears cogent to Samantha it doesn't follow that considering it will ineluctably lead her to believe that it is unlikely that the President is in New York City.[32] Cogency is a relative matter. After considering the media evidence and realizing that it indicates that the President is unlikely to be in New York City, Samantha would have to weigh it against what she sees as her clairvoyance evidence for the opposite conclusion. We can't rule out the possibility that Samantha's belief that the President is in New

York City emerges victorious from this contest, even if she conducts it honestly and conscientiously and with no goal other than truth. This possibility is not ruled out by BonJour's description of Samantha's case, and if this is the situation in which she finds herself, then according to ER2* her belief that the President is in New York City will be epistemically rational and responsible, from the point of view of her conception of the situation.

We have to conclude that BonJour's description of case 1 is compatible with situations in which Samantha's belief is epistemically rational and responsible, as well as with situations in which it isn't. However, there is no reason to doubt that BonJour would want to say in both situations that Samantha's belief lacks warrant. Hence, assuming that we require an explanation of the phenomenon that works in both situations, subjective irrationality is not the *explanans* we are looking for. Subjective irrationality can't explain why Samantha's belief lacks warrant because she might not be subjectively irrational in holding her unwarranted belief.

We can make the same point with respect to cases 2 and 3. The outcome is that the absence of warrant in cases 1–3 couldn't be explained in terms of subjective irrationality, even if warrant were incompatible with subjective irrationality (i.e., even if Premise 2 were true). This result undermines completely BonJour's defence of Premise 2. He argues that it is entailed by what he takes to be the best explanation of the absence of warrant in cases 1–3. But what he takes to be the best explanation of this phenomenon is no explanation at all, since its *explanans* is not even true in all instances of the phenomenon. Hence BonJour's argument for Premise 2 boils down to the claim that it is entailed by an explanation which doesn't work.

9. TWO FURTHER ARGUMENTS

I'd like to round up my discussion of BonJour's attack on reliabilism by considering two more aspects of his reasoning. The first one concerns BonJour's attempt to bolster his anti-externalist case with an analogy with ethics. In ethics we find the same contrast between objective and subjective perspectives for the assessment of action that BonJour exploits in the assessment of belief. BonJour starts by considering an externalist account of moral assessment, according to which the moral justification of an action is determined by its actual consequences, independently of the subject's beliefs. As BonJour

points out, saying that an agent's action is morally justified in this external sense "is not at all inconsistent with saying that his action was morally unjustified and reprehensible in light of his subjective conception of the likely consequences" (p. 44).

Then BonJour suggests that the moral externalist might try to accommodate this subjective sense of justification by adding to her account of when an action is morally justified "the further condition that the agent not believe or intend that it will lead to undesirable consequences" (ibid.). But BonJour finds this proposal unsatisfactory:

> Since it is also, of course, not required by moral externalism that he believe that the action will lead to the best consequences, the case we are now considering is one in which the agent acts in a way that will in fact produce the best overall consequences, but has himself *no belief at all* about the likely consequences of his action. But while such an agent is no doubt preferable to one who acts in the belief that his action will lead to undesirable consequences, surely he is not morally justified in what he does. (pp. 44–45)

For BonJour, the analogy with the epistemological case is perfectly clear:

> And similarly, the fact that a given sort of belief is objectively reliable, and hence that accepting it is in fact conducive to arriving at the truth, need not prevent our judging that the epistemic agent who accepts it without any inkling that this is the case violates his epistemic duty and is epistemically irresponsible and unjustified in doing so. (p. 45)

I want to make two points about this. First of all, notice that the problem in the moral case only arises on the assumption that on an externalist account of moral justification it cannot be required that the agent "believe that the action will lead to the best consequences". BonJour doesn't see the need to support this assumption, but presumably he thinks that an account of moral justification that accepted this requirement would no longer qualify as externalist. There is little point in quibbling about what makes an account qualify as externalist. However, the analogue of this position in epistemology is explicitly classified by BonJour as externalist. It is the position that he attributes to Dretske. It requires, BonJour tells us, "that the would-be knower also believe that the externalist condition is satisfied" (p. 233). Notice that, since it is not required that this belief has warrant, this position would still sustain a solution to the epistemic regress problem. I think there is little doubt that this position has nothing to fear from BonJour's ethical analogy.

The second point I want to make questions more generally the plausibility of the analogy. The problem that it raises for externalism in epistemology is that the externalist is committed to ascribing warrant to beliefs which the subject accepts "without an inkling" that they are "conducive to arriving at the truth". Presumably this notion is to be construed in analogy with the situation in ethics of an agent who "has himself *no belief at all* about the likely consequences of his action". In other words, the externalist is committed to ascribing warrant to beliefs that the subject is not inclined to regard as true. I want to suggest that the type of situation that this objection postulates simply doesn't arise. Believing a proposition is inextricably linked with the inclination to regard it as true. The externalist is not in danger of ascribing warrant to beliefs which aren't accompanied by this kind of inclination simply because there are no such beliefs.

Let me now turn to what is probably the key thought underlying BonJour's advocacy of the evidential constraint and his consequent rejection of externalism. Its starting point is a principle connecting epistemic justification and truth:

> The distinguishing characteristic of epistemic justification is [...] its essential or internal relation to the cognitive goal of truth. It follows that one's cognitive endeavors are epistemically justified only if and to the extent that they are aimed at this goal. (p. 8)

Notice that this principle by itself doesn't lend any support to the evidential constraint. In fact, it wouldn't be surprising if a reliabilist invoked it in support of her account of justification, claiming that aiming at the cognitive goal of truth requires using cognitive processes that are sufficiently likely to bring it about.

Of course one could object to this objectivist reading of "aiming at truth", on the grounds that it leaves the subject's conception out of the picture, and hence is not suitable for assessing the rationality and responsibility of her cognitive endeavours. A natural way of trying to bring that perspective into the picture would be to switch to a construal of aiming at truth in terms of the subject's beliefs and more basic inclinations concerning the extent to which her beliefs and cognitive processes bring her close to this goal. On this subjectivist construal, a subject would be aiming at truth so long as she didn't hold beliefs "against her better judgment", i.e., so long as she had done her best by her lights to determine the truth value of the propositions she believes. However, on this reading of "aiming at truth" BonJour's principle would still not lend support to the

evidential constraint. A reliabilist could accommodate the principle, on the subjective reading of aiming at truth, by adopting the proposal I made at the end of Section 7 to treat epistemic rationality and responsibility, construed along the lines of ER2*, as an additional necessary condition for warrant.

Thus neither on the objective nor the subjective reading of ''aiming at truth'' does BonJour's principle lend support to the evidential constraint. To use the principle in this way requires the specific reading of the notion that BonJour puts forward in the continuation of the passage quoted above, where he informs us that aiming at the cognitive goal of truth ''means very roughly that one accepts all and only those beliefs which one has good reason to think are true''. Needless to say, on this reading of ''aiming at truth'', BonJour's principle would support the evidential constraint. The problem for BonJour is that his reading stands itself in need of support, and I suspect that it is at least as vulnerable as the evidential constraint itself. For it is hard to see how someone who is inclined to reject the evidential constraint on warrant can be expected to be more open to being persuaded to accept an evidential constraint on aiming at truth. The evidential constraint on warrant can receive no effective support from this reading of aiming at truth, or a fortiori from BonJour's principle.

10. IRRELEVANCE

When a hypothesis is put forward as a solution to a problem, the proposal can in principle be resisted on two different fronts. On the one hand, one could argue that the problem can't be solved by the hypothesis because the hypothesis is false. On the other hand, one could argue that even if the hypothesis were true it would fail to solve the problem. The attempt to solve the epistemic regress problem with an externalist account of warrant is open to both kinds of attack. On the one hand, one could argue that externalism can't solve the problem because externalism is false – because it is wrong to construe warrant along externalist lines. On the other hand, one could argue that even if an externalist account of warrant were correct, the epistemic regress problem would remain unsolved – externalism would be, in this sense, irrelevant to the epistemic regress problem. Let me refer to this claim as the *irrelevance thesis*.[33]

As we have seen, BonJour's attack on an externalist solution to the epistemic regress problem seeks to undermine the proposal by

arguing that externalism is false – that it would be wrong to construe warrant along externalist lines. Notice, however, that BonJour explains his adoption of this line of attack as resulting from strategic considerations. It is his reaction to a difficulty generated by how radically externalism departs, in BonJour's view, from "the Western epistemological tradition" (p. 36):[34]

> [...] this very radicalism has the effect of insulating the externalist from any direct refutation: any attempt at such a refutation is almost certain to appeal to premises that a thoroughgoing externalist would not accept. (p. 37)

Strategic considerations aside, BonJour is convinced that the externalist solution to the epistemic regress problem should be rejected not only on the grounds that externalist accounts of warrant are false. Even if they were true they wouldn't solve the problem. This is the verdict that he attributes to Descartes and to "generations of philosophers who followed" (Ibid.): "the suggestion embodied in externalism would have been regarded as simply irrelevant to the main epistemological issue" (Ibid.). BonJour's support for this verdict is unambiguous:

> My own conviction is that this reaction is in fact correct, that externalism [...] reflects an inadequate appreciation of the problem at which it is aimed. (Ibid.)

It is hard to read this as anything but an endorsement of the irrelevance thesis.

However, I want to argue that BonJour would face a serious difficulty in trying to combine the irrelevance thesis with other aspects of his position. The difficulty arises from his support of a construal of the epistemic regress problem along the lines of what I have called the *contemporary regress argument*. The problem is simply that if externalism were correct, the problem generated by the contemporary regress argument would be solved, as the argument for the impossibility of knowledge would be deprived of a crucial premise.

BonJour might want to resist this outcome on the grounds that this purported solution to the problem generated by the contemporary regress argument involves "a repudiation of the very conception of epistemic justification or reasonableness as a requirement for knowledge" (p. 35), and hence that, instead of solving the problem, this position entails "that the regress problem in the form with which we are concerned would simply not arise, so that there would be no need for a solution [...]" (Ibid.). But BonJour doesn't want to extend this verdict to the versions of externalism that we have considered here.

And in any case it is hard to see what importance could be attached to the contrast between solving the problem and preventing it from arising, so long as the position under consideration has the consequence that the possibility of knowledge is no longer under threat.

The reason why BonJour can't argue for the irrelevance thesis is not the radicalism of the externalist proposal. The reason is rather that if the problem of the epistemic regress is construed as BonJour proposes, the irrelevance thesis is false. Perhaps the antecedent of this conditional is redundant, and the irrelevance thesis is false under any construal of the epistemic regress problem.[35] Then the problem would only afflict internalist accounts of warrant. But another possibility is that there is a different construal of the problem that renders the irrelevance thesis true. On such a construal, the problem would not be solved by the adoption of an externalist account of warrant. Whether such a construal of the problem can be provided is, I submit, still an open question – and one of the most important challenges for contemporary epistemology. BonJour might be right in thinking that such a construal exists, but he is certainly wrong in thinking that his own proposal fits the bill.

NOTES

* I have presented some of this material at UCL, Trinity College Dublin and the 10th Italo-Spanish meeting of Analytic Philosophy, held in Pamplona. I am grateful to these audiences.

[1] See L. BonJour, *The Structure of Empirical Knowledge*. Page references in brackets in the text are to this book. For BonJour's defence of coherentism, see ch. 5–8. For his attack on the doctrine of the given, see ch. 3.

[2] See L. BonJour, "The Dialectic of Foundationalism and Coherentism".

[3] Cf. Ibid., p. 117.

[4] Cf. Diogenes Laertius, *Lives of Eminent Philosophers*, IX 88–89, Sextus Empiricus, *Outlines of Pyrrhonism*, I 164. Aristotle appears to have been aware of the argument. Cf. Aristotle, *Posterior Analytics*, Chapter 3.

[5] BonJour's discussion of the problem is restricted to empirical knowledge. Here we shall raise the problem as affecting knowledge in general.

[6] BonJour formulates the problem in terms of the notion of *justification*. Its meaning is supposed to be given in independence of the connection with knowledge that I am treating as the definition of warrant, since BonJour thinks that there is room for positions according to which justification is not necessary for knowledge (cf. BonJour, *The Structure of Empirical Knowledge*, p. 35). However, he is not very explicit as to what's built into the notion. In any case, BonJour accepts the accounts of warrant with which we will be concerned as accounts of justification. Hence formulating the problem in terms of warrant won't distort BonJour's views for the purposes of this discussion. On this topic see Section 10.

[7] Proponents of the inferential account would want to weaken the requirement that the subject should have the belief connecting the target proposition with the evidence for it, while preserving the idea that the connection would have to be cognitively accessible to the subject in some way. The same point can be made concerning the subject's cognitive access to the propositions that play the role of evidence. These modifications wouldn't have a significant effect on the problems we are concerned with. Cf. Ibid., pp. 19–21 for BonJour's discussion of this issue. Notice also that one might want to demand some kind of positive epistemic status for the subject's belief in (or cognitive access to) this connection. However, the problem can be generated without this additional requirement.

[8] Inferential trees with infinite branches can be considered defective even if one accepts the possibility of having infinitely many beliefs. The issue is more complex than I'm making it sound, but I shall follow BonJour in leaving this possibility to one side. For a discussion of this point see E. Sosa, ''The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge'', Section 6.

[9] Notice that my claim is not that BonJour considers the argument sound, but that he regards it as the right construal of the problem.

[10] BonJour, ''The Dialectic of Foundationalism and Coherentism'', p. 132.

[11] D. M. Armstrong, *Belief, Truth and Knowledge*, p. 166.

[12] See A. I. Goldman, ''What Is Justified Belief?'', *Epistemology and Cognition,* ''Strong and Weak Justification''. For the contrast between the kind of view advocated by Armstrong and process reliabilism see F. F. Schmitt, ''Justification as Reliable Indication or Reliable Process?''

[13] The ideas presented in this chapter were first published in L. BonJour, ''Externalist Theories of Empirical Knowledge''.

[14] Cf. Also BonJour, *The Structure of Empirical Knowledge*, p. 50.

[15] Goldman has resisted this point. See below.

[16] Cf. BonJour, *The Structure of Empirical Knowledge*, pp. 41-42.

[17] At least in his initial reaction to BonJour's cases, Goldman seemed to concede the strength of this intuition. See below.

[18] Robert Fogelin has endorsed this diagnosis. Cf. R. J. Fogelin, *Pyrrhonian Reflections on Knowledge and Justification*, p. 45. Cf. Also J. Vogel, ''Reliabilism Leveled'', pp. 608.

[19] Cf. Goldman, *Epistemology and Cognition*, pp. 62-63. Notice that the non-undermining provision doesn't jeopardize the possibility of using reliabilism to underwrite the externalist strategy for dealing with the epistemic regress argument, since the resulting view still makes room for the possibility of warranted beliefs for which the subject lacks adequate evidence.

[20] Cf. Ibid., p. 112.

[21] Cf. A. I. Goldman, ''Epistemic Folkways and Scientific Epistemology''.

[22] See Ibid. for an alternative reliabilist diagnosis of clairvoyance intuitions.

[23] Using the intuitive approach to adjudicate the contest between reliabilism and the evidential constraint requires cases that neither party is likely to impugn as marred by contamination. We would need to focus on beliefs formed with a cognitive device exhibiting the following two features. On the one hand, its reliability should not be ruled out by fundamental aspects of our world-view. On the other hand, we should have no strong inclination to assume that the typical subject will have adequate evidence for or against its reliability. I have argued that Norman's case doesn't fit the

bill, but adequate cases are not hard to find. It seems to me that the intuitions that these cases elicit strongly support the reliabilist case, but I shall not pursue this point here.

[24] BonJour considers this objection in BonJour, *The Structure of Empirical Knowledge*, p. 46. For an early defence of this line of thought, see W. P. Alston, "What's Wrong with Immediate Knowledge?" The argument has been attacked by Richard Feldman (cf. R. Feldman, "Epistemic Obligations") and more recently by Matthias Steup (cf. M. Steup, *An Introduction to Contemporary Epistemology*).

[25] I am simplifying matters by thinking of the rationality and responsibility of an action or belief as a yes-or-no question. It would probably be more adequate to construe them as a matter of degree, but it should be clear how to adapt my formulations to achieve this. Notice also that the simplification introduces an implicit reliance on a notional threshold for the likelihood that A will promote G.

[26] This conception of epistemic rationality and responsibility is closely connected to Richard Foley's notion of egocentric rationality. Cf. R. Foley, *Working without a Net*. The only important respect in which the two notions differ from each other concerns the role that Foley accords to reflection. Thus, according to ER2* the beliefs of a subject who has done her best by her lights will be epistemically rational and responsible. The beliefs of such a subject will also be egocentrically rational provided that she trusts reflection – that she thinks that doing her best requires engaging in reflection. However, if the subject mistrusts reflection, her beliefs might be epistemically rational and responsible according to ER2* but not egocentrically rational. On this point, cf. Foley, *Working without a Net*, p. 97 and A. Plantinga, *Warrant: The Current Debate*, p. 133.

[27] On an adequate construal of the notion of doing one's best by one's lights to determine the truth value of *p*, one could fail to do so not only by failing to use a source of evidence that is related in this way to the proposition that *p*, but also by failing to use a source of evidence that is related in this way to the proposition that there is a source of evidence that is related in this way to *p*… A simple induction would take care of all levels of iteration.

[28] The text actually suggests a different line of reasoning, according to which what follows from Norman's lack of adequate evidence for his belief is not that it is subjectively irrational, but that it lacks warrant. However, in the context of trying to support the evidential constraint on warrant, this move would be utterly question begging.

[29] For the possibility of subscribing to the evidential constraint on warrant but not to the parallel constraint on epistemic rationality and responsibility, cf. Susan Haack's position on epistemological justification and moral justification: "I do not think it is always morally wrong to believe on inadequate evidence. […] however, I think it is always epistemically wrong to believe on inadequate evidence – in the sense that believing on inadequate evidence is always epistemologically unjustified belief". S. Haack, "'The Ethics of Belief' Reconsidered", p. 138. Bruce Russell also makes room for this possibility. He writes: "questions of epistemic blamelessness are separate from questions of the goodness of evidence". B. Russell, "Epistemic and Moral Duty", p. 38.

[30] Plantinga, *Warrant: The Current Debate*, p. 45.

[31] Ibid.

[32] If this connection held, then one could argue for the irrationality of Samantha's belief without invoking the simplifying assumption. Then doing her best by her lights would put her in the curious situation of believing that the President is in New York City and believing that this belief is unlikely to be true, which, according to ER2, would render her first-order belief epistemically irrational.
[33] Prominent proponents of the irrelevance thesis include Barry Stroud and William Alston. cf. W. P. Alston, *The Reliability of Sense Perception,* B. Stroud, ''Understanding Human Knowledge in General''. See E. Sosa, ''Philosophical Scepticism and Epistemic Circularity'' for an attack on Stroud's position, and R. Fumerton, *Metaepistemology and Skepticism* for an attack on the compatibility of Alston's endorsement of the irrelevance thesis with his commitment to externalism.
[34] For a different conception of this tradition, cf. F. F. Schmitt, *Knowledge and Belief*.
[35] Cf. M. Bergmann, ''Externalism and Skepticism'' for a recent defence of this claim.

REFERENCES

Alston, W. P.: 1983, 'What's Wrong with Immediate Knowledge?' *Synthese* **55**, 73–95.
Alston, W. P.: 1993, *The Reliability of Sense Perception*, Cornell University Press, Ithaca, N.Y.
Aristotle: 1994, *Posterior Analytics,* 2nd ed. Clarendon Press, Oxford.
Armstrong, D. M.:1973, *Belief, Truth and Knowledge*, Cambridge University Press, London.
Bergmann, M.: 2000, 'Externalism and Skepticism'. *Philosophical Review* **109**, 159–94.
BonJour, L.: 1980, 'Externalist Theories of Empirical Knowledge.' in P. French, T. Uehling, Jr. and H. K. Wettstein (eds.), *Midwest Studies in Philosophy,* Vol. V: *Studies in Epistemology*, University of Minnesota Press, Minneapolis. pp. 53–73.
BonJour, L.: 1985, *The Structure of Empirical Knowledge*, Harvard University Press, Cambridge, Massachusetts.
BonJour, L.: 1999, 'The Dialectic of Foundationalism and Coherentism.' in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology*, Blackwell, Malden, Massachusetts and Oxford.
Diogenes Laertius.: 1925, *Lives of Eminent Philosophers*, William Heinemann, London.
Feldman, R.:1988, 'Epistemic Obligations.' in J. E. Tomberlin (ed.), *Philosophical Perspectives, 2, Epistemology*, Ridgeview, Atascadero, California. pp. 235–56.
Fogelin, R. J.: 1994, *Pyrrhonian Reflections on Knowledge and Justification,* Oxford University Press, New York and Oxford.
Foley, R.: 1993, *Working Without a Net,* Oxford University Press, New York.
Fumerton, R.: 1995, *Metaepistemology and Skepticism,* Rowman & Littlefield, Lanham, Maryland.
Goldman, A. I.: 1979, 'What Is Justified Belief?' in G. S. Pappas (ed.), *Justification and Knowledge*, Reidel, Dordrecht. pp. 1–23.

Goldman, A. I.: 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
Goldman, A. I.: 1988, 'Strong and Weak Justification.' in J. E. Tomberlin (ed.), *Philosophical Perspectives, 2, Epistemology*, Ridgeview, Atascadero, California. pp. 51–69.
Goldman, A. I.: 1991, 'Epistemic Folkways and Scientific Epistemology.' in *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, MIT Press, Cambridge, Massachusetts and London, England, 155–75.
Haack, S.: 1997, '"The Ethics of Belief" Reconsidered.' in L. E. Hahn (ed.), *The Philosophy of Roderick M. Chisholm*, Open Court, Chicago and La Salle, Illinois. pp. 129–44.
Plantinga, A.: 1993, *Warrant: The Current Debate*, Oxford University Press, New York and Oxford.
Russell, B.: 2001, 'Epistemic and Moral Duty.' in M. Steup (ed.), *Knowledge, Truth, and Duty*, Oxford University Press, Oxford and New York. pp. 34–48.
Schmitt, F. F.: 1981, 'Justification as Reliable Indication or Reliable Process?' *Philosophical Studies* **40**, 409–417.
Schmitt, F. F.: 1992, *Knowledge and Belief,* Routledge, London.
Sextus Empiricus.: 1990, *Outlines of Pyrrhonism*. R. G. Bury (ed.), Prometheus Books, Buffalo, New York.
Sosa, E.: 1980, 'The Raft and the Pyramid: Coherence Versus Foundations in the Theory of Knowledge' in P. French, T. Uehling, Jr. and H. K. Wettstein (eds.) *Midwest Studies in Philosophy,* Vol. V: *Studies in Epistemology*, University of Minnesota Press, Minneapolis. pp. 3–25.
Sosa, E.: 1994, 'Philosophical Scepticism and Epistemic Circularity,' *Aristotelian Society Supplementary Volume* **68**, 263–90.
Steup, M.: 1996, *An Introduction to Contemporary Epistemology,* Prentice Hall, Upper Saddle Rever, N.J.
Stroud, B.: 1989, 'Understanding Human Knowledge in General' in M. Clay and K. Lehrer (eds.), *Knowledge and Scepticism*, Westview Press, Boulder, Colorado.
Vogel J.: 2000, 'Reliabilism Leveled,' *Journal of Philosophy* **97**, 602–623.

Philosophy Department
Univeristy College London
Jower Street
London WC1E 6BT
England
email: j.zalabardo@ucl.ac.uk