

## Scepticism, Externalism and Predictive Dimension of Knowledge Claims

ŽIVAN LAZOVIĆ

Department of Philosophy, University of Belgrade, Čika Ljubina 18–20, 11 000 Belgrade, Serbia  
zlazovic@f.bg.ac.rs

ORIGINAL SCIENTIFIC ARTICLE / RECEIVED: 07–07–11 ACCEPTED: 01–08–11

---

ABSTRACT: Ordinary knowledge claims are challenged by philosophical scepticism which holds that we are unable to exclude the possibilities of error involved in well-known sceptical alternatives (e.g., the brain-in-a-vat hypothesis). In order to explain how we can resist this challenge, first I compare philosophical and ordinary doubt. I point out that they do not differ in terms of the way they aim to undermine knowledge claims, but rather *in the character of the alternatives to which they appeal*. Thus, in ordinary contexts, philosophical sceptical alternatives should not be considered relevant because they are farfetched, not supported by any indication that they really might exist in the given circumstances. Since this point concerns the assertability rather than the truth of our knowledge claims, I further argue that while our evidential basis for their assertability is internal, their truth depends upon certain assumptions concerning the causal history of our beliefs, reliability of our cognitive abilities, success in identifying and excluding relevant alternatives, etc. In everyday knowledge attributions, these assumptions operate as *externalist* preconditions so that we may know something only if they are correct even without knowing that they are correct. Finally, I point out an implicit predictive dimension of our knowledge claims consistent with their fallibility: when we properly claim to know something, we do not imply impossibility of being mistaken, but rather hold that we are *not* wrong and that no sudden twist of future events will show us to be wrong.

KEY WORDS: Externalism, knowledge, sceptical alternatives, scepticism.

---

### Introduction

Before we encounter a philosophical sceptic, we have confidence in a number of propositions about the external world – at least those propositions for which we have strong and convincing evidence. Thus now, while I am working on this paper, I am ready to claim without any hesitation that

I know that I am typing on my computer, that I live in Belgrade, or, after watching the NBA Eastern Conference final series on TV, that the Miami Heat eliminated the Chicago Bulls. However, once we become acquainted with the problem of philosophical scepticism, we get an impression that such everyday knowledge claims are shaky and tend to be less willing to simply accept them. In her attempt to undermine what we count as knowledge, the philosophical sceptic typically employs the following strategy. She relies on a seemingly plausible *general knowledge requirement* (GKR) according to which, in order to know that  $p$ , we must be able to rule out at least those possibilities which we take to be incompatible with our knowing that  $p$ . The most obvious are possibilities of error. Suppose that on the base of empirical evidence  $E$  a person  $S$  claims to know that  $p$ . Relying on GKR, the sceptic devises and puts forward hypotheses which are such that (a) they specify possible situations in which  $\neg p$ ; (b) by their very nature, they are consistent with  $E$  and cannot be ruled out by any piece of  $E$  that  $S$  might want to use; and (c) as incompatible with the truth of  $p$ , if true, they prevent  $S$  from knowing that  $p$ . According to GKR, since such hypotheses invoke possibilities of error that are inconsistent with  $S$ 's knowledge claim, it seems that  $S$  has to rule them out in order to have knowledge that  $p$ .

We will refer to hypothetical situations incompatible with our knowing that  $p$  as *alternatives*. Sceptical alternatives provide at least partly causal explanations of our beliefs because they suppose that a belief's causal history is radically different from the way we think it to be. The well-known dreaming hypothesis is such an alternative. While forming beliefs about the external world, we have only our perceptual experience to appeal to. We believe that, in normal circumstances, ordinary physical things around us trigger our perceptions through which we acquire knowledge of the external world. However, all our perceptual experience may be equally well explained by the hypothesis that we are simply dreaming the things we claim to know, and since nothing can reveal to us whether we are awake or not, the world around us might be radically different from the way it looks to us. Thus, the philosophical sceptic concludes that we lack knowledge about the external world, since we cannot exclude the possibility that we are dreaming and that the things around us are completely different from the way we took them to be.

An easy way to construe sceptical alternatives relies on the highly intuitive *epistemic closure principle* (ECP) which states that knowledge (and possibly justification) is closed under entailment: if  $S$  knows that  $p$  and knows that  $p$  entails  $q$ , then  $S$  may get to know that  $q$  by competently inferring it from  $p$ .<sup>1</sup> The dreaming hypothesis is usually presented in this

---

<sup>1</sup> The proviso "by competently inferring" is needed in the standard formulation of the principle, since it might happen that someone knows that  $p$  and that  $p$  implies that  $q$  without inferring that  $q$ . See Williamson (2000: 117); Hawthorne (2004: 32).

way: I claim to know that I am sitting by the desk and typing on my computer. I presumably also know that if I am sitting by the desk and typing on my computer, I am not asleep in my bed. According to ECP, I should know that I am not asleep in my bed. But since my perceptual experience in dreams might look the way it looks when I am awake, how can I exclude the possibility that now, while believing that I am sitting by the desk and typing on my computer, I am asleep in my bed merely dreaming that I am sitting by the desk and typing? Another example is the brain-in-a-vat hypothesis: I claim to know that I am typing on my computer. I presumably also know that if I am typing on my computer, I am not a bodiless brain in a vat, victim of a malicious scientist who supplies me with false experiences that I am typing on my computer. If I were being manipulated in such a way, my experiences in the brain-in-a-vat situation might be exactly the same as in the actual circumstances, so that I cannot rule out the possibility that I am not deceived. The sceptic concludes that since I do not know that I am not a brain in a vat, I do not know that I am typing on my computer. The brain-in-a-vat alternative undermines, through the ECP, my supposed knowledge because its truth is inconsistent with the truth of the proposition I claim to know; obviously, I cannot be typing on my computer and also be a bodiless brain in a vat with all my experiences supplied by a malicious scientist.

Employing this strategy, the philosophical sceptic may accomplish her goal through a number of slightly different ways. It seems that any plausible definition of knowledge should produce a list of necessary conditions which are jointly sufficient for having knowledge. Take, for instance, the familiar *justified true belief* (JTB) definition of knowledge according to which S knows that  $p$  if and only if  $p$  is true, S believes that  $p$ , and S has a *justified* belief that  $p$ .<sup>2</sup> Any sceptical alternative will be effective if it undermines at least one of three necessary conditions listed in the JTB definition. This leaves room for a variety of sceptical alternatives whose knowledge-destroying capacity does not directly stem from ECP and their obvious inconsistency with the truth condition.

The dreaming hypothesis and the brain-in-a-vat scenario are sketched here as possibilities which include that  $\neg p$ , but they also can be construed in a somewhat different way; the dreaming hypothesis may suppose that our dreams come true, and the brain-in-a-vat scenario may include a benevolent scientist who deliberately supplies our brains with true perceptual experiences and memories. Despite ensuring the truth of  $p$ , these sceptical alternatives are still incompatible with our knowledge that  $p$ : even if my

---

<sup>2</sup> I do not want to defend JTB analysis; the same line of reasoning that follows would apply whatever definition is offered.

belief that the Heat eliminated the Bulls is true, I do not know that if this belief, together with all background evidence I have for it, is merely a content of my dreaming or part of my brain-in-a-vat experience. So construed, these possibilities clearly threaten our knowledge by attacking the justification and not the truth condition, because the truth of my true belief that the Heat eliminated the Bulls in these hypothetical scenarios appears as a lucky guess.

It should be noted that, on some additional assumptions, even the belief condition might be a target of sceptical alternatives. If the so-called thesis of semantic externalism is correct, the content of at least some of our beliefs is partly determined by their causal origin in external objects.<sup>3</sup> For example, someone who has never been causally connected in the proper (direct or indirect) way with the Heat and the Bulls as NBA basketball teams cannot form beliefs about them. So, the content of my belief that the Heat eliminated the Bulls is partly determined by its actual causal history, leading from the real Eastern Conference finals games through my watching the series on TV, or reading NBA news, and culminating in my belief that the Heat eliminated the Bulls. At least in the radical brain-in-a-vat scenario, in which I am always supposed to be a brain in a vat, the content of my beliefs is not causally connected with real things in the proper way, so that what I believe in this scenario could not be the belief about *the Heat's winning*; it seems that, comparing with my actual belief about the Heat's winning, my brain-in-a-vat belief would be identical qualitatively but differ in content – something like a belief about the Heat's winning as a mental image contained in my brain-in-a-vat experience.<sup>4</sup> Still, according to this scenario, I would (falsely) believe that my brain-in-a-vat belief is a belief about the real Heat's winning; in other words, I would mistakenly identify the content of my belief expressed by the sentence “The Heat eliminated the Bulls”. Thus, given the externalist view on mental content, the philosophical sceptic might use some of her favourite alternatives to argue that we are not able to satisfy even the belief condition for knowledge about the external world. From a subjective point of view, our brain-in-a-vat beliefs will look to us phenomenologically the same as in normal circumstances. Due to the fact that we are ignorant about their

<sup>3</sup> The thesis has been defended in Putnam (1999). Others have tried to develop a semantic anti-sceptical argument that proceed from externalist thesis to the conclusion that one is not a brain in a vat (see Brueckner 1999). Here I do not want to deal with this kind of response to the classical problem of scepticism with respect to the external world; rather, I would like to point out how semantic externalism construed as an account of “the mind/world relationship” (Putnam 1999: 31) gives a rise to a second-order scepticism which Brueckner calls “*skepticism about knowledge of content*” (Brueckner 1999: 48–9).

<sup>4</sup> See Putnam (1999: 35).

causal origins, and as we cannot exclude the possibility that what looks to us as the belief that  $p$  is, after all, just a belief that  $p$ -in-the-mental-image, then we can never know that  $p$ .<sup>5</sup>

The challenge of the philosophical sceptic lies in her demand that, to defend ordinary knowledge claims, we must rule out her alternatives by proving them false. There is a variety of responses to this challenge, from heroic attempts to show that we can somehow rule out sceptical hypotheses, to rejecting the ECP or restricting the GKR to the domain of the so-called relevant alternatives.<sup>6</sup> My aim is not to provide a definitive answer to philosophical scepticism. Rather, in this paper, I will follow the line of the linguistic response, inspired by J. L. Austin, which holds that in everyday situations, we rightly disregard the hypotheses put forward by the philosophical sceptic since they fail to satisfy some normative requirements for reasonable doubt in our ordinary epistemic practice. In the light of Stroud's objection that Austin's account deals with the assertability but not the truth of knowledge claims, I will argue that an Austinian line of argument is promising only if we pay attention to the externalist dimension of knowledge displayed in some assumptions implicitly operating in our knowledge attributions. The content and the role of these requirements for reasonable doubt and externalist assumptions will be shown in more detail through a comparison of ordinary and philosophical sceptical strategies. Finally, I will try to improve our epistemic prospect of warding off the sceptic's attack by uncovering an implicit predictive dimension in our ordinary knowledge claims.

### Ordinary and Philosophical Doubt

The conflict between our ordinary knowledge claims of the form "S knows that  $p$ " (where  $p$  stands for a contingent proposition) and the conclusion reached by the philosophical sceptic is striking. There is no dispute about the main sceptical diagnosis of our epistemic position; all our knowledge claims of this kind are fallible and, as such, always vulnerable to doubt. The question is whether the typical philosophical doubt which appeals to possibilities like those in dreaming and brain-in-a-vat hypotheses is strong enough to undermine ordinary knowledge claims. The proponents

---

<sup>5</sup> In Lazović (2000) I have incorporated sceptical alternatives, depending on the target condition in JTB analysis, into the truth-alternatives, belief-alternatives, and justification-alternatives to our knowledge claims.

<sup>6</sup> For example, Putnamian anti-sceptical strategy is heroic in the way Descartes's reply to the evil genius hypothesis was, Dretske (2001) adopts the relevant alternatives theory without the ECP, while Stine (1976) and Cohen (1988) present contextualist versions of the relevant alternatives theory which retain the ECP.

of linguistic anti-scepticism point out that the form of doubt on which the philosophical sceptic builds her arguments imposes higher standards than does our ordinary epistemic practice.<sup>7</sup> They charge the philosophical sceptic with distorting or misunderstanding the meaning of the ordinary concept of knowledge. If they are correct, their diagnosis provides firm grounds for a plausible explanation of why, in everyday life, we ignore the philosophical sceptical alternatives and why they do not constitute a genuine challenge to our ordinary knowledge claims. In other words, the doubt raised by these alternatives is robbed of strength because the alternatives do not respond to core requirements that apply in everyday epistemic practice which are closely related to the ordinary concept of knowledge. After some initial worry over the sceptic's conclusion, the realisation that it originates in a lack of understanding or in a skewed interpretation of the ordinary concept of knowledge brings a sigh of relief.

Doubt is not restricted to philosophers, however. It is frequently exhibited in everyday epistemic practice, both in quotidian life and science. I claim to know that the Heat eliminated the Bulls because I read in the paper that the Heat beat them last night. In such an ordinary context, I am not disturbed by dreaming and brain-in-a-vat alternatives, but I may have overlooked that the Heat now leads three games to one and needs one more victory to win the series. Unlike dreaming and brain-in-a-vat alternatives, this possibility of error raises a doubt strong enough to force me to withdraw my knowledge claim.

What are the key differences, then, between ordinary and philosophical doubt? When we compare some examples of the two, we discover that they do not differ in their *argumentative structure* (they both rely on GKR) but in the *character of the alternatives* to which they appeal in their attempts to undermine knowledge.

Suppose that I am visiting the National Museum in Belgrade when my attention is drawn to Renoir's well-known painting, "The Bather." In the JTB definition of knowledge, in order for me to know that I have exactly that painting before me, I have to satisfy three conditions: I should believe it is Renoir's "Bather," my belief needs to be true, and my belief should be justified. On the one hand, from the traditional, internalist point of view, the fulfilment of the belief condition looks least problematic, because the internalist typically presupposes that in normal circumstances one has some sort of privileged access to the content of her experiences, and it seems there is no way in which she could possibly be deceived in believing that *p*. On the other hand, the truth condition conveys the conceptual point that in searching for knowledge, one strives to ascertain whether her

---

<sup>7</sup> Stroud (1984: 40) calls this anti-sceptical strategy diagnostic.

beliefs are true or not. In ordinary contexts, truth is assumed to be an objective relation between our beliefs (or, more precisely, propositions which express the content of our beliefs) and the states of affairs in the external world. Hence, the truth is, from the outset, introduced as an *externalist* condition, that is, as a condition whose fulfilment should be indicated by, but does not depend on, our epistemic position. The justification condition enters the stage at this point. Its role, traditionally construed, is to bridge the gap between the subjective and objective ingredients of knowledge; it is expected that justification will direct beliefs towards the truth so that the truth of our justified beliefs is not a matter of chance.<sup>8</sup>

Arguably, in everyday situations like visiting museums, all three conditions can be met easily; while looking at it, I believe the painting is Renoir's "Bather", this belief is true, and, what is more, it is justified. The available evidence could include the fact that I am in a museum, I have good lighting to examine the picture, and I am well acquainted with the artist's work (recognisable technique, characteristic subject, brushstrokes, colours, chromatic range, layers, etc.). Even those who are far from expert should be able to read the text below the painting where the name of the artist and the title of the work are usually displayed. How in circumstances like these could anyone question my knowledge claim that the painting I am looking at is Renoir's "Bather"?

We are already familiar with the sceptic's general strategy. Bearing in mind the actual situation and the evidence *E* in the light of which *S* believes that *p* is true, the sceptic will put forward a hypothetical scenario *H* that is consistent with *E* but which, if actualised, would prevent *S* from satisfying at least one of the necessary conditions for knowing that *p*. In this way, granting GKR, the burden of proof is placed firmly on *S*'s shoulders: if the possible truth of *H* is inconsistent with *S*'s knowing that *p*, it seems that *S* is epistemically committed to show that *H* is not true.

How would such a strategy of doubt creation work in our example? It may be helpful to recall events at a local modern art gallery about twenty years ago. The curator of this gallery, a skilled master of painting techniques, hatched a scheme to become rich. He made very good copies of the works of several well-known national painters, hung the forgeries on the walls of the gallery, and sold the originals on the black market. For years, the public enjoyed the work of their modern national artists without noticing they were fakes. The fraud only came to light when the friend of

---

<sup>8</sup> Here it should be noted that, traditionally, justification also has been spelled out in *internalist* terms. A typical justification for a belief is understood as the evidence cognitively available to the person having the belief and conditioned by the person's cognitive perspective.

a collector noticed a painting in the collector's private collection which was almost identical to work on display in the gallery. Sceptics adore such cases. While during the visit to the National Museum in Belgrade, I view a painting which, based on appearances, I justifiably believe to be a Renoir, the sceptic easily envisages a possibility similar to the one happened in the local modern art gallery. Perhaps someone has cunningly forged the original, removed it, and hung the forgery in its place. The scenario may be richly embroidered with details, becoming more credible and realistic: perhaps the swap has been made for the same criminal reasons as in the gallery case, possibly there are security concerns, bearing in mind that "The Bather" was once stolen, or maybe the original requires restoration and the director of the museum does not want visitors to have to forgo the pleasure of "seeing" the painting. These details are not terribly important to us at this juncture. What is important is that this hypothetical scenario undermines my supposed knowledge that I am looking at the original painting.

The hypothesis that the painting is a forgery is obviously inconsistent with my statement that painting before me is Renoir's "Bather" and attacks, *via* the ECP principle, the truth condition of my knowledge claim. How can I know in the given circumstances that this painting is Renoir's "Bather"? Since the sceptic's hypothesis includes the assumption that the forgery is nearly identical to the original, it seems that I cannot rule out this possibility by any piece of my initial evidence. However, I may not have to admit defeat. The GKR urges me, in the light of the possibility of error hypothesised by the sceptic, and which I, in the given circumstances, did not consider, to refrain from claiming to know unless I can rule out that possibility. But the content of this hypothesis is such that I am (at least in principle) in a position to check it, and depending on the result, I can accept or reject the sceptic's objection. At my disposal are some (as Austin would note)<sup>9</sup> "recognised" techniques and procedures for addressing such questions. For example, I could approach the director of the museum, as a likely reliable "source" of information, or speak to experts in a laboratory capable of establishing whether the work is original.

In short, the sceptic's forgery hypothesis, as presented to me, is only a *temporary* challenge to my efforts to reach knowledge. It suggests that, in the given circumstances and in the light of the initial evidence, I do not know that I am looking at the original painting. It does not suggest that it is impossible for me to establish that I am. Possible sceptical hypotheses of this kind are countless, but they all point to neglected obstacles or defects in my initial evidence which could, once revealed, be removed. They only

---

<sup>9</sup>Austin (1970: 87).



confirm Dretske's lesson<sup>10</sup> that the road to knowledge is hard, requiring great epistemic effort and careful thought as to the methods used to form justified beliefs.

The situation becomes much more complicated when the *philosophical* sceptic appears on the scene. The alternatives she raises aim to shake the foundations of our epistemic aspirations and lead us into an epistemic dead end. Her goal is not to show that, in certain circumstances, we have not arrived at knowledge because we have carelessly neglected or overlooked certain possibilities that threaten our knowledge claim. Rather, she aims to show that we can never know anything because possibilities exist that undermine all our knowledge, and, however hard we try, we will never be in a position to rule them out. More specifically, one could, in the aforementioned situation, observe the following: since it is possible for us to undergo experiences in dreams that are qualitatively the same as those in our waking life, at the moment when I believe myself to be in the National Museum, looking at Renoir's "Bather", it is *possible* that I am only dreaming of this. Alternatively, it is possible that all I experience at that moment is simply the creation of a powerful deceiver who has the ability to manipulate the content of my mind. These alternatives seem to involve possibilities which point to unbridgeable chasms in our epistemic path. Whatever test we try in our attempt to prove they are not actualised (that we are not dreaming and not unfortunate victims of a deception), that effort can easily be assimilated into the above-mentioned hypothetical scenarios. These possibilities apparently undermine my knowledge claim in the same way as the possibility that the painting is a fake. If it were true that I am dreaming or that someone had created within me the impression that I am in the National Museum, then I certainly could not claim to *know* I am looking at Renoir's "Bather", as the justification condition or the truth condition of our knowledge claim would not be satisfied. However, unlike the previous situation, these sceptical alternatives do not simply elicit temporary challenges to my particular knowledge claim in the given circumstances. They easily spread over almost all other accompanying knowledge claims, because within the dreaming or brain-in-a-vat scenario I would not know neither that I am looking at Renoir's "Bather", nor that I am standing in front of painting, nor that I am in the National Museum, nor even that the external world exists. In contrast to the possibility that the painting I am looking at is fake these alternatives reach deeper, threatening the very basis of my epistemic aspirations and allowing the philosophical sceptic to conclude triumphantly that I (or anyone else) can never know anything.

---

<sup>10</sup> As Dretske notices, "Believing is easy, knowledge is hard" (Dretske 1983: 3).

As we have noted above, ordinary and philosophical doubt share the same argumentative structure. If I claim to know that I have before me Renoir's "Bather", and the ordinary sceptic suggests that it might be a clever forgery, while the philosophical sceptic suggests that perhaps I am dreaming or I am being systematically deceived by someone who is intervening in our minds, my knowledge claim is being challenged in the same way. In both cases, I am faced with possibilities which are incompatible with my supposed knowledge, and, according to the GKR, I have to rule them out if I wish to defend my knowledge claim. In spite of the fact that they are sharing the same strategy, the aims with which ordinary people and philosophers express their doubts are very different, as are the character of the alternatives they construe and the conclusions they reach.<sup>11</sup>

Firstly, ordinary doubt is *local* in the sense that it is usually *directed at a specific single belief* or, at most, to such a belief *along with a small group of closely related or implied beliefs and apart from all other beliefs*. In our example of Renoir's "Bather", the question "how do you know that it is really Renoir's 'Bather', as it might be a cunning forgery" is directed only at my belief that I am looking at the original one. The question in no way challenges my background knowledge consisting of other justified true beliefs, that I am living in Belgrade, that I am visiting the National Museum, that I am looking at a painting, etc. Moreover, the suggestion that I do not know it is Renoir's "Bather" until I rule out the possibility that it is a cunning forgery does not imply that I could not, in other circumstances, get the knowledge which, in the given circumstances and at the given moment, is being undermined. In contrast, as it is pointed out in our example, the hypotheses of the philosophical sceptic tend to be *global, involving alternatives which threaten to undermine all (or virtually all) beliefs we hold in the given circumstances*. If I am dreaming, or if I am systematically deceived, it appears that not only can I not know I am looking at Renoir's "Bather", but I cannot know I am living in Belgrade, visiting the National Museum, looking at a painting, etc.

Secondly, as directed at individual beliefs of whatever kind, ordinary doubt *does not tend towards generalisation*. The alternatives which it puts forward are not such as to allow the doubt to spread out from the given belief to the overarching set of beliefs to which it belongs (for example, from one particular perceptual belief to all perceptual beliefs, or from one belief at which we have arrived by inference to all such beliefs of this kind). Unlike this, even when the philosophical sceptic is dealing with a particular belief, she takes it as a representative for the set of beliefs to

---

<sup>11</sup> I have provided more detailed discussion of similarities and dissimilarities between ordinary and philosophical doubt in Lazović (2000).

which it belongs, *striving to spread her doubt to all beliefs of the given kind* (for instance, all perceptual beliefs, all beliefs concerning past or future, all beliefs about the external world, and so on).

Finally, as we have seen in the example of Renoir's "Bather", ordinary sceptical alternatives only encourage us to strengthen our epistemic position by presenting temporary incidental challenges to our knowledge claims. In contrast, the philosophical sceptic is looking for alternatives which are incompatible with almost everything we are supposed to know but are so cleverly designed that it seems impossible to unearth any evidence that they are not true. Relying on the GKR that demands us to rule out such alternatives if we wish to keep our knowledge claims, this line of reasoning leads to the *radical* sceptical conclusion that we *can never* know anything.

Can we somehow respond to the philosophical sceptic and protect our ordinary knowledge claims from such a radical and disastrous conclusion?

### The Linguistic Approach

As we have seen, ordinary and philosophical doubt do not differ in the way they threaten our knowledge claims. Rather, due to divergent aspirations, their difference concerns the content and characteristics of the alternatives through which they are expressed. The linguistic response maintains that the philosophical sceptic, by using her alternatives to undermine our ordinary knowledge claims, in fact, raises the bar for the standards of knowledge so high that she ends up distorting the ordinary meaning of the verb "to know". At first glance, this diagnosis may offer relief; if the philosophical sceptic's alternatives invoke requirements that ordinary use of the verb "to know" does not have to meet, that is, if they imply a concept of knowledge which is different from the ordinary one, the sceptic's conclusion, however radical, will not conflict with our everyday knowledge claims.

An earlier version of the linguistic approach<sup>12</sup> points out that ordinary epistemic practice imposes requirements of rationality regarding beliefs that qualify for knowledge, as well as requirements as to the strength of doubt. In other words, if we demand that beliefs be justified, we are expecting doubt to be reasonable. This suggests that GKR should be modified so that the domain of sceptical alternatives which are able to endanger knowledge in everyday situations should be restricted to those possibilities which are grounded on a reasonable doubt, i.e. restricted to the scope

---

<sup>12</sup> This version can be found in Austin (1970: 84).

of the so-called relevant alternatives. During the 1970's and 1980's, such a general thesis was developed (by Dretske and Nozick, for instance<sup>13</sup>) into a more specific view which focuses on the denial of ECP. While the philosophical sceptic interprets GKR in the spirit of infallibilism, understood as a view that *all conceivable* alternatives incompatible with one's knowledge must be ruled out, the relevant alternatives view holds that in everyday epistemic practice we are expected to be in a position to rule out only the *relevant* alternatives.<sup>14</sup> The main issue is, of course, under which conditions could a doubt be deemed reasonable, i.e. what makes an alternative relevant?

Recall Austin's description of our epistemic practice. To arrive at specific knowledge, we employ different, well established techniques and procedures, appropriate to a particular type of the case at stake.<sup>15</sup> Regarding the situation of identifying a painting in a museum, such contextually adjusted methods would include reliance on one's prior acquaintance with a given artist's work, reading the identification label, seeking an expert's opinion, etc. The application of these procedures, to use Austin's terminology, puts us in a position to know about the painting. Of course, everyday epistemic practice allows reservations elicited by ordinary doubt. This is why the conditions in which we claim to know something clearly take into account a set of raised as well as anticipated objections which bring that knowledge claim into doubt. So in response to the question, "how do you know that painting *A* is by the artist *N*?" we would normally reply by a) showing that we are really in a position to know because of our prior knowledge, the suggestive context, the identification procedures carried out, and so on, and/or b) drawing attention to the particular features of the object (the age of the canvas, recognisable technique, choice of subject, brush strokes, colouring, chromatic scale, layers of paint, etc.) on the basis of which we have identified it. At this point, the sceptic could threaten our knowledge claim by a) questioning the reasons on the basis of which we assert that we have that particular painting before us ("Look closer – don't you see that, in fact, it does not have the property *F* which you thought it had", when possession of *F* is a necessary condition for it to be the painting we claim it to be), or b) questioning whether the reasons we have given

<sup>13</sup> See Dretske (2001) and Nozick (1981). Some of Austin's remarks could easily be worked up into an abandonment of ECP (cf. Austin 1970: 83–84).

<sup>14</sup> The relevant alternatives view is typically spelt out in terms of "ruling out". Recent contextualist approaches (DeRose 1992, 1995; Lewis 1999; Cohen 2000) hold that the criteria of relevance are dependent on the knowledge attributors in the given context. A persuasive critique of the epistemic contextualism as a semantic thesis can be found in Hawthorne (2004) and Stanley (2005).

<sup>15</sup> Austin (1970: 87).

in support of our knowledge claim are sufficient for knowledge (“Isn’t it the case that there is another painting with the feature *F*, and you should rule out the possibility that this is not, in fact, painting *B*, or a copy of *A*”, and so on). Both in quotidian life and science, ordinary doubt employs objections of this type.

Philosophical doubt appeals to alternatives of a different type. While ordinary doubt is usually provoked by some *particular* defect in our epistemic position which might be a source of error, philosophical doubt is based on quite general and farfetched considerations about logical possibilities which seem so far removed from the real circumstances that we usually do not feel compelled to consider them seriously. The ordinary sceptic would respond with: “That isn’t sufficient because the painting could still be a forgery”. In contrast, the philosophical sceptic would turn to one of her favourite hypotheses, saying, “That isn’t enough, because you may be a brain in a vat, being supplied by the false experience that you are in a museum and have the painting before you”.

Austin had in mind two requirements for reasonable doubt.<sup>16</sup> The first bears on the *specificity* of a doubt. If someone objects, saying that we do not know this is painting *A* by artist *N*, because all the evidence we have provided to that effect is not sufficient for us to know that fact, she must indicate what is wrong with our epistemic position or tell us what is missing from our evidence. In other words, she is obliged to specify a hypothesis that can explain how we arrived at the (erroneous) belief that this is painting *A* by artist *N*, how we arrived at our initial evidence, and in what respect that evidence is inadequate.

The second requirement pertains to the *relevance* of a doubt. Even a specific alternative will not be taken seriously unless the sceptic confronts us with a specific reason for believing that this alternative *might obtain* in the given circumstances. Since we have based our claim that we know that *p* on evidence *E* which consists of reasons *for* believing that *p*, we are right to expect that the sceptic should provide reasons *against* this belief and make her doubt reasonable. At the very least, she should point out characteristics of the situation which suggest that, in spite of all evidence we got, her hypothesis might be true or specific indications, in light of which her alternative does not appear to be a mere logical possibility but, rather, something that could really have happened, thereby allowing us to be wrongly convinced that we know.

Philosophical doubt usually does not breach the specificity requirement. In our example, the brain-in-a-vat hypothesis specifies the way in which our evidence may be insufficient, as much as does the hypothesis

---

<sup>16</sup> Cf. Austin (1970: 84).

that the painting is a forgery. However, the philosophical sceptic does breach the relevancy requirement because she mainly relies on generalised and abstract logical possibilities without pointing to any specific reason to believe that her hypothesis might be true. That explains why, according to the linguistic approach, her doubt does not represent a genuine challenge to our knowledge and why we ignore her alternatives in everyday situations. Unless it is shown that the doubt is both well specified and relevant, we are not obliged to take it seriously and rule out the alternatives involved in it. Thus, the burden of proof is shifted to the philosophical sceptic who would have to show not only that there are possible alternative explanations of how we get to the evidence, but that we failed to notice some specific indication that her hypothesis might be true in the given circumstances. And if she decides to employ this manoeuvre, she confronts a dialectical trap: in order to make her alternative relevant, she needs to *specify certain characteristics of the situation which indicate that it might be true*; but she will thereby step onto the terrain of ordinary doubt where we may, having been informed of what we failed to take into account and after investing further epistemic effort, test and possibly confirm or exclude her reasons for doubt. Having in mind that her aim was to undermine almost all ordinary knowledge claims, it turns out that her alternatives cannot be both effective and relevant at the same time.

Hence the philosophical sceptic wrongly assumes that *any* alternative which is incompatible with our knowledge claims endangers our knowledge. In everyday epistemic practice, we are faced with more realistic demands; when we claim to know something we are not obliged to rule out *all* alternatives – only those which, due to the circumstances, are deemed relevant, that is, those for which it can be shown that there are specific reasons to believe that they might be true. Because the philosophical sceptic invents alternatives which are specified but not relevant, the conclusions to which she comes are not in genuine conflict with our ordinary knowledge claims. Only ordinary doubt remains legitimate, appealing to alternatives which meet both the requirements of specificity and relevancy but which, exactly for this reason, could be (at least in principle) checked.

Barry Stroud has outlined a possible reaction for the philosophical sceptic.<sup>17</sup> According to Stroud, the linguistic approach oversimplifies matters in so far as it supposes that the issue is terminological. He points out that with knowledge claims we can, as in the case of any statement, distinguish between their *assertability* and their *truth* conditions, that is, the conditions under which a statement can justifiably be asserted, and the conditions under which that statement is true. In the light of this distinc-

---

<sup>17</sup> Stroud (1984), particularly ch. 2.

tion, Stroud concedes that Austin's thesis addresses only the assertability and not the truth conditions of ordinary knowledge claims. While with respect to the first of these, there is indeed no conflict, with respect to the second, there is enough room for philosophical scepticism. In a particular situation, the ability to apply all the recognised procedures and rule out all relevant alternatives adds up to having the right to claim that we know something. But having the right to claim to know something does not guarantee that the truth conditions for this claim have been met, that is, that we really know. Building on a well-known article by Thompson Clark,<sup>18</sup> Stroud observes that in quotidian situations governed by practical needs and concerns, we can make justified knowledge claims and, at the same time, fail to meet the conditions of knowledge. Of course, these practical constraints affect the domain of alternatives that can be considered relevant. The alternatives that the philosophical sceptic suggests seem too distant and abstract to be considered relevant for practical purposes. However, those possibilities might be considered relevant for meeting the truth conditions of knowledge claims. The philosophical sceptic may still allow for the difference between "knowing for practical purposes" and "really to know"<sup>19</sup> stressing that, just because in everyday circumstances, we feel no obligation to deal with possibilities such as the dream or the brain-in-a-vat deceptions, this does not mean that we have really obtained the knowledge we were looking for.

Recall the example borrowed from Clark, of airplane spotters, trained (for some important practical purpose) to distinguish planes of type *A* from those of type *B*. The training manual states that type *A* planes have the features *FGH*, while type *B* planes have the features *FGI*. But it turns out that there is another type of plane, *C*, which, at least when it is observed from the ground, cannot be differentiated from type *B*, because it too has the features *FGI*. This type is, however, rare, and distinguishing it from *B* is considered, at least for the purposes of observation, not important and is not mentioned in the manual. In such circumstances, it seems fitting that the careful spotter who spots a plane with the features *FGI* claims to know it to be a *B* type. Following the rules in the manual and not knowing that the type *C* plane exists, the spotter will be justified in claiming that she knows it to be a type *B* plane. Nevertheless, because from time to time, a

---

<sup>18</sup> Clark (1972).

<sup>19</sup> Contextualists hold that meaning and truth value of knowledge claims depend upon practical interests of knowledge attributors. Stanley also maintains the thesis that whether *S* knows that *p* is partly determined by *S*'s practical interests, but unlike contextualists he defends it as a "metaphysical thesis about the nature of knowledge relation" (Stanley 2005: 120). Both views amount to the claim that knowledge is always tied to practical purposes.

type *C* may appear, the spotter will not know that it is a *B*, even if it is a *B*, unless and until she excludes the possibility that it is a type *C*.

Stroud's point is that, in general terms, our epistemic position is similar to that of the airplane spotter. Whenever we have a right to claim to know something, it may turn out that we are wrong because we have neglected an alternative incompatible with our knowledge, thereby failing to meet the truth conditions of our knowledge claim.

The analogy, however, cannot be stretched far enough to support philosophical doubt. The alternative that Stroud employs in this example is much like that used in ordinary doubt, because (although the spotter may not be aware of it) a plane of type *C* exists. Yet the mere logical possibility that such a plane exists should not bother the spotter. If the sceptic confronts the spotter with such a hypothesis, in order to make her doubt reasonable, not only she will have to specify her hypothesis, but she must point to an indication which suggests that her hypothesis might be true in the given circumstances. Otherwise, no reason can be seen why the spotter should consider the possibility of the presence of plane type *C* rather than the presence of any further plane type –  $D_1$ , or  $D_2$ , ..., or  $D_n$ , for example, whose features *FGI* are similar to those of type *B* but differ from it in some additional features. The status of a sceptical alternative and, consequently, the strength of a doubt changes substantially depending on its relevance. The philosophical sceptic is still confronted with a dilemma: in order to achieve her aim and find an irremovable hypothesis which admits the generalisation of doubt, she must turn to an alternative that is only specified and not relevant; to really threaten knowledge, she would have to make the suggested alternative relevant, thus returning it to the domain of ordinary doubt and making it testable.

However, the philosophical sceptic still might insist that ECP commits us to excluding at least those alternatives presented to us, whose truth is inconsistent with the truth of the proposition we are claiming to know. As noted before, some philosophers have tried to explain how knowledge is possible without ECP taken in a unrestricted sense. Yet, I stick to the majority reluctant to reject such highly intuitive principle. The ECP belongs to the most recognised and established procedures for obtaining knowledge, namely, procedures of deductive inference. And there is at least one sense in which it can be used even with respect to the alternatives raised by the philosophical sceptic. Since ECP is based on *modus ponens* as a truth-preserving principle, if I rightly claim to know that *p* ("I have hands"), and I am acquainted with the philosophical sceptic's alternative that *q* ("I am the brain in a vat"), the ECP enables me to conclude that not-*q* and justifies me in claiming to know that not-*q*. Furthermore, if in the



given circumstances I do know that  $p$  (which by definition means that  $p$  is true), and I am faced with the alternative that  $q$ , I can get to know that not- $q$  by employing the truth-preserving ECP, i.e., by competently deducing not- $q$  from  $p$ . And if there is no other, independent method of getting to know that not- $q$ , as apparently is the case with the “brain in a vat” alternative, the ECP remains the only available way to get to the knowledge of not- $q$  (provided, of course, that I rightly claimed to know that  $p$ , and  $p$  is true). The philosophical sceptic still might complain that what was really questionable from the beginning was my knowing that  $p$ . First of all, however, my way of getting to know that  $p$  has included other procedures (perception, for instance) for acquiring knowledge; it is not obvious that I have to use the same method for getting to know the logical consequences of  $p$ . Secondly and, as we will see in the next section, more importantly, *our way of getting to know that  $p$  involves some externalist assumptions concerning the causal structure of the world and the reliability of employed procedures* which to some extent guarantees the truth of  $p$ ; if we rightly claimed to know that  $p$  and all externalist conditions (including the truth condition) for our knowing that  $p$  are satisfied, then there is nothing wrong in the employment of ECP as a reliable way of acquiring knowledge of any logical consequence of  $p$ , including not- $q$ . Finally, in the next section we will also point out how the aforementioned externalist assumptions are reflected in *the predictive dimension* of our knowledge claims.<sup>20</sup>

### Externalist Assumptions and Predictive Dimension of Knowledge Claims

In the framework of the JTB conception of knowledge, Stroud’s distinction between the assertability conditions and the truth conditions of knowledge claims mirrors a gap between the justification condition, traditionally construed in an internalist manner, and the externalist truth condition. This gap creates room for hypothetical scenarios in which our beliefs, however we may justify them, turn out to be wrong (as in some of the sceptic’s scenarios) or, if true, true only by coincidence (as in well known Gettier-like examples). That is the reason why externalists about knowledge maintain that the link between the belief and the truth condition in the JTB definition, even when it is mediated by the justification, needs to meet some further, externalist requirement.<sup>21</sup> The core externalist intuition is that the

<sup>20</sup> The predictive success provides perhaps the strongest case for our ordinary knowledge about the external world.

<sup>21</sup> Externalism about knowledge need not be inconsistent with internalism about justification, i.e. it can allow that some knowledge involves internalist justification, denying

subject's belief, apart from being true and justified, has to be connected to external facts in some nomic or modal way in order to constitute knowledge. This intuition was best captured by Nozick's idea of tracking the truth of the proposition believed, and further elaborated either as the sensitivity (were  $p$  false,  $S$  would not have believe that  $p$ ) or the safety requirement ( $S$  would believe that  $p$  only if  $p$  is true).<sup>22</sup> The details of externalist epistemologies need not concern us here. For the purposes of this paper it suffices to stress that, while internalist view of justification may account our practice of making knowledge claims, externalism accommodates the intuition that knowing is more than just having a right to say "I know"; when I rightly claim to know, I (reasonably) expect that my justified belief *is* true, but the belief's success in tracking the truth finally depends on the causal structure of the external world.

My primary goal in this section is to show how the basic externalist idea is manifested in some assumptions involved in our practice of making ordinary knowledge claims. For instance, when we claim to know that  $p$ , we are not only convinced that we have successfully ruled out all relevant alternatives to  $p$  and showed that  $p$  is true, but we are also assuming that the causal history of our belief that  $p$  is not radically different from the way we take it to be and the method we used for getting knowledge was reliable. Furthermore, our ordinary knowledge claims typically rest on some additional assumptions which also play an externalist role, because we are used to take them for granted without really knowing them.<sup>23</sup> When we make a knowledge claim, we tacitly assume plenty of things without supposing that we know them independently. When I say that I know I am typing on my computer, I take for granted, among many other things, that everything is normal and no strange scenarios like being a brain in a vat are going on. Attributing knowledge to someone else, others likewise assume the same sorts of things; under the circumstances, also assuming that the circumstances are normal, my wife and my son would agree that I know I am typing on my computer. What we are assuming while making the knowledge claim belongs to the externalist dimension of knowledge; when we take these assumptions for granted and claim we know that  $p$ , our knowledge claim is true only if the assumptions are correct, even if at that time we do not know they are correct.

---

only that all knowledge is internalist. The minimal externalist thesis is that knowledge is at least partially determined by some external facts which the subject needs not to be in position to know.

<sup>22</sup> The sensitivity requirement has been proposed by Nozick (1981), and the safety requirement was introduced by Sosa (1999) and defended by Williamson (2000: ch. 8).

<sup>23</sup> In ordinary circumstances of our knowledge claims, most of these assumptions may be rendered as pragmatic or conversational presuppositions.

Some of the things assumed are such that we do not know them at the time of making a knowledge claim, but we may (at least in principle) arrive at a position to know them. Visitors to a museum typically take it for granted that the paintings displayed there are not cunning forgeries. When they are claiming that Renoir's "Bather" is before them, they simply take for granted that it is not a forgery. They can know that it is Renoir's "Bather" without knowing at that time that it is not a forgery; they know this only if their initial assumption is correct – even if they do not know at that time that it is correct. If anyone specifies an alternative scenario in which the painting is a forgery and indicates a compelling reason for this alternative in the given circumstances, thereby making it relevant, the visitors would be not only in a position, but also epistemically obliged, to discover whether it is the original painting or a forgery if they wish to defend the claim that they know it to be Renoir's "Bather". As shown in the section on ordinary and philosophical doubt, this is the essential dialectic of the debate between knowledge attributors and ordinary doubt sceptics.

The same dialectic is at work in the airplane spotter example. Supposing that planes of type *C* exist and appear in the sky from time to time, the alternative that when the spotter says the plane she is observing is a *B*, the possibility that it is a *C* might be relevant. If there are some compelling reasons why planes of the type *C* are not mentioned in the training manual (for instance, they exist but few are left and those few are no longer in use), and if there are no indications that a plane of type *C* is appearing in the sky, the spotter has a right to make her knowledge claim. Moreover, the spotter may know that the plane she is observing is a *B* without knowing at that time that it is not a *C*; she knows this only if her implicit assumptions that the manual and her vision are reliable in the given situation and that the plane she is observing is not some *C*-like plane which cannot be discerned from a *B* plane with the help of the manual are correct. The situation changes if someone objects that planes of type *C* appear in the sky from time to time and points to indications (for instance, those who wrote the manual overlooked the fact that planes of type *C* exist and fly across the area)<sup>24</sup> that the plane the spotter is observing at that moment might be a *C* instead of a *B*. Then, the alternative that it is a plane of type *C* instead of a *B* becomes relevant, and the spotter will not have the right to make her knowledge claim until she rules out this alternative. As a farfetched possibility, the mere hypothesis that there are *C*-like planes belongs to the domain of philosophical doubt, and since it does not meet the relevancy

<sup>24</sup> Of course, there are some complexities here. As Dretske and some contextualists have shown, the stringency of the relevancy condition may vary depending on knowledge attributors' epistemic standards, interests or some other contextual features.

condition, it cannot threaten the spotter's ordinary knowledge claim. By making it relevant, we step onto the terrain of ordinary doubt where the spotter has a way, at least in principle, of checking out the doubt, so that she can either rule it out or withdraw her knowledge claim.

Among the things we typically take for granted in everyday epistemic practice, there are some we can never know independently: for instance, that we are not brains in vats. When, in the light of all available evidence, I rightly claim to know that I am typing on my computer, I take it for granted that I am not a brain in a vat without supposing that I know I am not a brain in a vat; and I know I am typing on my computer only if the assumption that I am not a brain in a vat is correct, even if I do not know that it is correct. The philosophical sceptic is right in claiming that I cannot know that I am not a brain in a vat, at least not in the way I get to know most things in the external world, including the fact that I am sitting beside my desk and typing on my computer.<sup>25</sup> But the assumption that I am not a brain in a vat functions like an externalist precondition for my knowing other things about the external world; if it is correct, then I do know that I am typing on my computer. Moreover, as illustrated at the end of the last section, it seems there is one way in which I can know I am not a brain in a vat: if I become faced with the brain-in-a-vat alternative and realise that my typing on the computer entails that I am not a brain in a vat, I can competently use ECP and claim that, *in this way only*, I know that I am not a brain in a vat. Although the typical justification for our knowledge claims is internalist, in making these claims we rely on some externalist assumptions concerning the external world and our cognitive position. Due to the externalist dimension of knowledge, I may know that *p* without knowing that I know that *p*. I may know that I am typing on my computer even if I do not know that all preconditions for my knowing are satisfied; and if I rightly claim and, in the externalist sense, do know that I am typing on my computer, even if I cannot know that I am not a brain in a vat in the same way (i.e. through perceptual evidence) in which I can know that I am typing on my computer, it might be allowed that I may know this by competently using ECP.<sup>26</sup>

<sup>25</sup> As noted, the philosophical sceptic may specify her alternatives, but she can hardly make them relevant. The dialectic of the debate demands that an alternative, to be considered, should be relevant, and it seems that the philosophical sceptic's alternatives can do their job only if they are *not* relevant.

<sup>26</sup> Basically, this is a neo-Moorean maneuver. I agree with Pritchard that neo-Mooreanism as an externalist position has advantages over its two main rivals, the relevant alternatives view and contextualism, since it keeps the ECP and avoids some implausible relativistic consequences of contextualism. I also agree that the prospects for an epistemologically *internalist* answer to philosophical scepticism are dim (cf. Pritchard 2005: 67–8).

Furthermore, the bundle of externalist assumptions concerning the nature and structure of the world external to us and our epistemic position within the causal nexus of events is reflected in a tacit predictive dimension of our knowledge claims.<sup>27</sup> When we make a knowledge claim, we do not imply that it is not possible to be mistaken; rather, we are implying that we *are not* mistaken. This involves predictive expectations without a commitment to infallibility. When I am claiming now to know that I am typing on my computer, relying on the evidence which is at my disposal, I am completely sure that I am not wrong, but I do not want to claim that I know that there is no possibility that I am wrong. Based on the background of externalist assumptions that accompany my knowledge claim, I expect that the world or the causal history of my belief is not radically different from the way we are usually supposed to be: no funny things like the brain-in-a-vat scenario are going on, the computer I am typing on will not suddenly disappear from my desk or turn into a heap of plastic pieces, and so on.<sup>28</sup> Even though it is logically possible that outrageous events such as these happen at any instant, in the everyday circumstances in which we make our ordinary knowledge claims, there are no indications that they are really happening, and most of the time they do not.

Stroud is right in holding that knowing something is more than just having a right to make a knowledge claim. However, seeing what this “more” is, involves bringing more externalism into our conception of knowledge, i.e. realising that our ordinary knowledge claims are accompanied by many externalist assumptions concerning past, present, and future courses of events. As noted before, when we rightly claim to know that *p*, the truth of our knowledge claim partly depends upon the truth of these assumptions, although we do not claim that we know these assumptions to be true. Hence, even when it seems that we have a right to make them, and even when we imply that we are not wrong, our knowledge claims may be defeated in the sense that in the near or far future, some strange twist of events may arise and show we were wrong. The predictive dimension of knowledge claims is consistent with general fallibility if we

---

<sup>27</sup> Austin has maintained that the phrase “I know” is *not predictive* “in such a way that the future can always prove it wrong” (Austin 1970: 88–9). He has in mind that an “outrageous” course of events may show that we *were not* right to say “We know” even though, at the time we were making the knowledge claim, we *had a right* to do so (*ibid.* 98). As we will see, my contention that knowledge claims *are predictive* is consistent with this fallibilist intuition and concerns only the fact that, when I claim to know that *p*, I claim that *p* is true.

<sup>28</sup> For the most part, these expectations are grounded on the ECP; i.e. they concern the truth of all propositions about the future which are entailed by the proposition we claim to know. They are also of practical importance, since we take our knowledge claims as a basis for planning the future, making decisions and acting.

interpret it in the externalist sense: when we (rightly) claim to know that  $p$ , among other things we are assuming we are also implicitly taking for granted that the future course of events will not prove us wrong; whether we know that  $p$  partly depends on the truth of this assumption, whether we know it or not.

As we can see, the general fallibility of our rightly asserted knowledge claims is not at issue. Yet the logical possibility of being refuted does not stack up against the possibility of knowing. The philosophical sceptic is right in holding that for all we know, we may turn out to be mistaken, but she is wrong to conclude that we never know. Of course, there have been cases in which we turned out to be wrong, and there will be cases in which we will turn out to be wrong again. In spite of that, and as far as we know, in even more cases, we turned out to be right, and there will probably be more instances in which we will turn out to be right again. Having epistemically done our best, we can still claim to know, relax, and let nature do the rest.

### References

- Austin, J.L. 1970. "Other minds", in *Philosophical Papers* (Oxford: Oxford University Press), 76–116.
- Brueckner, J. 1999. "Semantic answers to skepticism", in DeRose and Warfield (1999: 43–60).
- Clark, T. 1972. "The legacy of scepticism", *The Journal of Philosophy* 69, 754–769.
- Cohen, S. 1988. "How to be a fallibilist", *Philosophical Perspectives* 2, *Epistemology*, 91–123.
- . 2000. "Contextualism and skepticism", *Philosophical Issues* 10, 94–107.
- DeRose, K. 1992. "Contextualism and knowledge attributions", *Philosophy and Phenomenological Research* 52, 913–929.
- . 1995. "Solving the skeptical problem", *Philosophical Review* 104, 1–52 (repr. in DeRose and Warfield 1999: 183–220).
- . and T.A. Warfield (eds.). 1999. *Skepticism: A Contemporary Reader* (New York, Oxford: Oxford University Press).
- Dretske, F. 1983. "The epistemology of belief", *Synthese* 55, 3–19.
- . 2001. *Perception, Knowledge and Belief* (Cambridge: Cambridge University Press).
- Hawthorne, J. 2004. *Knowledge and Lotteries* (Oxford: Clarendon Press).

- Lazović, Ž. 2000. "Između običnog i filozofskog skepticizma", *Filozofski godišnjak* 13, 149–170.
- Lewis, D. 1999. "Elusive knowledge", in DeRose and Warfield (1999: 220–242).
- Nozick, R. 1981. *Philosophical Explanations* (Oxford: Oxford University Press).
- Putnam, H. 1999. "Brains in a vat", in DeRose and Warfield (1999: 27–42).
- Pritchard, D. 2005. *Epistemic Luck* (Oxford: Clarendon Press).
- Sosa, E. 1999. "How to defeat opposition to Moore", *Philosophical Perspectives* 13, 141–153.
- Stine, G.C. 1976. "Skepticism, relevant alternatives, and deductive closure", *Philosophical Studies* 29, 249–261.
- Stroud, B. 1984. *The Significance of Philosophical Scepticism* (Oxford: Clarendon Press).
- Stanley, J. 2005. *Knowledge and Practical Interests* (Oxford: Clarendon Press).
- Williamson, T. 2000. *Knowledge and its Limits* (Oxford: Oxford University Press).