

João Carlos Salles

ERNEST SOSA'S TELIC VIRTUE EPISTEMOLOGY¹

Abstract: Ernest Sosa's work recently took the form of a theory of telic normativity. As he presents, telic normativity is inherent to actions and attempts that characterize human performances, being telic because they are aimed at ends and often normative because we say they are better if successful and, therefore, if they reach their objective. The specific aims of this paper are, at first, to show the core of his new theory, notably his model for evaluating epistemic performances, associated with the shifting of position in his work on epistemic modalities such as "sensitivity," "safety," and "security." In a second moment, we suggest how Sosa is going forward just after achieving the current stage of his reflection – that is, formulating a "dawning light epistemology."

Keywords: Ernest Sosa, normativity, epistemic modalities, telic virtue epistemology, dawning light epistemology.

1. Ernest Sosa is one of the most important contemporary philosophers. His intellectual contribution has been remarkable on the philosophical scene for around 60 years. Furthermore, it is a unique contribution. After all, since 1974, his work constitutes its proper field of reflection, the virtue epistemology, recently taking, in 2021, the form of a theory of telic normativity, which now, as a process in full swing, deepens as a Dawning Light Epistemology.

In this text, we will seek to present the unique model of normative evaluation of human performances, particularly epistemic performances, taking into account his 2021 book *Epistemic Explanations: A Theory of Telic Normativity and What It Explains*. This model shows us how Sosa, on the one hand, always faced the challenges of the research program sparked by Edmund Gettier, not accepting *tout court* the reprimand to the program made by Timothy Williamson.

1 This research is supported by the Brazilian agencies CAPES and CNPq.

To this end, we will briefly present (1) the characteristic features of the traditional definition of knowledge as justified true belief, (2) Gettier's famous critique, and (3) Williamson's objection. Such a succinct presentation highlights Sosa's path, offering him only a context. Thus, on the other hand, we will show how Sosa developed a complex explanation for a multifaceted phenomenon such as knowledge without limiting himself to a definition that would supposedly resist successive counterexamples. He successfully established the normative frameworks of a taxonomy.

Given the scope of this text, we will not be able to address specific points we nevertheless mentioned. Sosa's reflection is, after all, a constant work in progress. Moreover, even the recent publication of this normative model only propels him to the new steps that he has already presented at several conferences.

In this way, the specific aim of this presentation is to show a quintessential aspect of his 2021 theory (notably, the normative model and the shifting of position in his work on epistemic modalities such as "sensitivity," "safety," and "security"). However, it is also essential to indicate that Sosa is going forward just after achieving the current stage of his reflection – that is, now formulating a "dawning light epistemology," with which he benefits a lot from his more recent and peculiar dialogue with Wittgenstein's *On Certainty*.

2. Knowing is not as trivial as it seems. The term has different meanings, and it is often crucial to reduce it so that the epistemic aspect stands out. Let us take knowing here as an action focused on facts, whose truth would be presented through declarative sentences. This is the meaning that interests Sosa throughout his career. As a result, pragmatic components, however important they may be and despite the centrality they increasingly acquire in Sosa's work, are subordinated to semantic aspects, being decisive, for example, the relationship between the notion of knowledge and notions such as evidence, belief, and justification. Knowing would be, then, for Sosa, equivalent to getting it right through statements.

Some conditions are primarily accepted, as stated by the traditional view of knowledge as Justified True Belief (JTB). (i) For something to be known, it must occur and can then be represented in a true proposition. So, for something to be known, it must be true. There may even be knowledge of what it is not, but precisely that it is not. The reflection is extensive

in the history of philosophy, as in the Aristotelian lesson that a proposition tells the truth if it says things are what they are or if it says they are not what they are not – a lesson, moreover, reiterated by Sosa: “What is not so cannot be known to be so: if anyone is to know that p then at a minimum it must be so, it must be true, that p .” (Sosa, 1994, p. xi)

Of course, the truth of p is thus a necessary but insufficient condition for knowledge. After all (ii), much of what is true goes wholly unnoticed and, therefore, unknown. We must take this into account and believe that something is (or is not) in order to know about it. A subjective act is required, manifesting as a belief or a judgment, whether or not it can be proven. Therefore, to believe that p is also a condition for us to know p , but obviously, it is still not enough.

However, (iii) belief alone does not bring the truth. Furthermore, even the conjunction between belief and truth can be a coincidence, as when we get a multiple-choice question correct from a mere guess. In addition to believing that something is accurate, we need a justification for our belief – with which knowledge would be a true and justified opinion. As said in *Theaetetus*, in a classic version of the traditional definition, the possibility of presenting a path to truth is also an indispensable feature of knowledge: “(...) when someone gets hold of the true judgment of something without an account, his mind is in a state of truth about it but doesn't know it.” (Plato, 2014, pp. 95–96 (202c))

3. Suppose we were satisfied with this view of knowledge as Justified True Belief. In that case, we should only ask for more perfect justifications since merely to believe being justified is not objectively being justified – as believing in following a rule is not following that rule, to remember a Wittgensteinian dictum.

However, even that favorable JTB conjunction, although required, would be insufficient. With the famous Edmund Gettier's 1963 paper, “Is Justified True Belief Knowledge?”, we learned better: the conjunction between belief, truth, and justification does not guarantee that, finally, we know. Gettier's problem, with precision and clarity, would reveal the scandalous inadequacy of the standard definition of knowledge.

The two counterexamples of his paper became immediate classics, one conveying logical consequence in a calculus structure of analyzed propositions and the other in a calculus structure of unanalyzed propositions. In both cases, long story short, the epistemic consequence does not follow the logical consequence, and so we may arrive at the truth by chance. Therefore, the subjective search for more objective justifications would not satisfy our epistemic claims, as even a justification according to logical inferential guidelines may fall short of knowledge.

Taking the simple Gettieresque example of inference by addition, we have that, from the truth of 'p', we can infer the truth of the disjunction between this proposition and any other since the truth of 'p' is enough to guarantee the truth of 'p v q'. The logical consequence is perfect. It is indifferent whether 'q' is true or false, being sufficient the truth of 'p'. However, the proposition 'p' being false, the truth of the disjunction is reached by mere chance, thus depending on the truth of 'q', which was simply unknown before. Our belief about the truth of 'p v q' would be accurate and justified, but it would not be a case of knowledge since it does not follow the pattern of the intended normativity that should surround our getting it right.²

4. Because of such undesirable possibilities raised by Gettieresque situations, a rich and specific research program has been developed with great technical sophistication, trying to answer this central question: What are the necessary and jointly sufficient conditions for us to have any knowledge and, therefore, for our eventual access to truth not to be a work of chance?

Indeed, a logical justification only sometimes carries with it epistemic correctness. Whoever seems to arrive reasonably at the truth may still be threatened by illicit favors of fortune. So, from the 1960s onwards, we have a community prepared to use heavy analytical artillery to operate the minor technical aspects – an extensive community equipped with the appropriate tools to get down to the nitty-gritty of what previously did not even seem like a problem.

Much ink was thrown into papers, re-editing (we must admit) a kind of scholasticism, now with analytical philosophy's rigor and scientific flavor. So, it seemed to some to be a relief to imagine that such a research program could cease.

After all, despite the results and valuable side products, the program awakened by Gettier came to be seen as perhaps idle or stimulated by an illusion. When someone thinks they have reached a safe place, they likely will face the truth that there are many ways of being wrong.

If someone, on the contrary, is justified, it would be expected that the consequences drawn from initial beliefs do not reach the truth by mere

2 Sosa claimed that the Gettier problem was what first gripped him in epistemology. Indeed, he published one of the first reactions to the Gettier paper in 1964 and summarized its two examples: "Suppose S has good evidence for his belief that p, from which in turn he deduces that p v q. But, unknown to S, (~p) & q. So, all three conditions for knowledge specified in the view under examination are fulfilled; but we still do not want to say that *S* knows that p v q." (Sosa, 1991, p. 15)

chance. In the case of luck, the beliefs being well founded, some inferential standard was disrespected. If so, the definition of knowledge would be in order. Although vague, it would be sufficient, even without indicating under what conditions it would be satisfied.

Nevertheless, being so, the definition itself would not rule out the possibility of a false justification, nor would it finally indicate what a reasonable justification would be, just as the competence of archers would manifest itself in how they can hit the target, even though they may miss.

5. It is beyond our intention and almost impossible to recompose the research program's very twists of examples and counterexamples. We only intend to highlight one milestone in the program's evolution, Ernest Sosa's virtue epistemology,³ which offers, for instance, a sufficient response to one of the strongest objections to the program itself, the criticism of Timothy Williamson,⁴ who believes that it is not possible to submit the notion of knowledge to an analysis, seeming to return the notion to the condition of a truism, without which would not even be possible to think and which, therefore, we fail to analyze.

Despite the strong influence of Williamson, with his formula of great rhetorical appeal (knowledge first!), the program remains alive, although now more for its stability than for its effervescence. After all, Williamson would only be right if our task were to look for definitions in the form of necessary biconditionals.

Moreover, the race would be vain and idle if we were hunting for counterexamples in which the equivalence between a justified true belief and knowledge would be false. In this case, above all, the notion of knowledge would be present in some way, even in a non-obvious way, in the examples of justified true belief, whose description would supposedly rule out such a possibility:

3 Facing the challenges posed by the analysis of knowledge, Sosa developed a path all his own with the proposition of a virtue epistemology, which is markedly normative. This truly original elaboration has its inception, we believe, in 1974, with the text "How do you know?". The famous 1980 text, "The Raft and the Pyramid," is also recognized as a seminal document of its genesis, so Sosa's epistemology of virtues can have two beginnings, both deserving of wide celebration. Regarding these two origins and the characteristic features of Sosa's epistemology of virtues, see our paper Salles, J. C. "The Amazing Mr. Magoo: To celebrate Virtue Epistemology's 50th anniversary," published in Spanish, with the translation by Carlos Caorsi, as "El increíble Sr. Magoo: Para celebrarel 50 aniversario de la Epistemología de la virtud," *Revista Elenkhos*, Vol. 6, N. 2, December 2023.

4 Timothy Williamson's *Knowledge and Its Limits* is deservedly one of the most influential contemporary epistemology books.

So, the interesting philosophical analysis of knowledge would go beyond the mere correlation of K with JTB, even beyond the proposed *necessary* correlation of K with JTB. The interesting philosophical proposal would be rather of the following form:

K is present, when it is, in virtue of, or grounded in, JTB. Or: Always, when JTB is present, then K is *thereby* present. And, moreover, when K is present, that is *because* JTB is present. (Sosa, 2023a, p. 2)

Even admitting that Williamson was correct in many aspects, Sosa then states: “epistemologists interested in Gettierology were addressing interesting philosophical explanations of how knowledge comes to be, of how it is metaphysically grounded.” (Sosa, 2023a, p. 2) In short, we are not in the game of mere analysis, but in that of philosophical explanation – a game that shall never cease and always needs justification. In other words, the normativity of the epistemic endeavor sustains its proper relevance and continuity – something that Sosa’s theory of normativity beautifully exemplifies.

6. Let us show a few traits of Sosa’s theory of telic normativity.

First, his defense of such an explanatory exercise (never dissociated from analysis) implies preserving one specific aspect of the philosophical activity of analyzing knowledge, namely, its normative dimension tantamount to understanding knowledge as a unique example of human action aimed at purposes – which is, moreover, better exemplified by alethic purposes. Therefore, Sosa’s theory of general human performances reiterates essential traits of a normative perspective when applied to knowledge.

Of course, normativity has several meanings, but, in any case, if we think about epistemic normativity, it is appropriate to emphasize that truth has a normative tint. Remembering Anselm’s lesson, something is true when it is as it ought to be. Thus, in the end, the truth is a matter of correctness, the correctness appropriate in each instance. A proposition normatively must be able to mean what it intends to express with it, just as it must be able to express, when declarative, how things are; that is, it must be able to saywhathappens correctly. (Cf. King, 2006, pp. 214–219)

So, we can use the notion of normativity in epistemology’s field in a somewhat limited but quite reasonable way, indicating with it the select set of procedures adopted when we aim to achieve the truth competently and thus obtain knowledge. Thus, “epistemic norms” are simply the procedures we follow when thinking and reasoning competently.

Second, Sosa’s model for evaluating performances makes evident, being normative, the central phenomena of a telic theory, such as attempt, success, competence, aptitude, and achievement:

If an archer shoots at a certain target, we can assess that shot in various respects. First, does it *succeed*? Does it hit the target? Second, how competent is the shot? The arrow may exit the bow with an orientation and speed that would normally take it straight to the bullseye. Even if a gust diverts it, the shot might still be competent. It can be *adroit* without being *accurate*. And it can be accurate by luck, without being adroit. But even a shot that is both accurate and adroit might still underperform. An arrow adroitly released from a bow may be headed straight to the bullseye when a gust diverts it so that it would now miss the target narrowly, except that a second gust eases it back on course. The archer succeeds in that attempt to hit the target, and the shot is also competent, as the arrow leaves the bow perfectly directed and with the right speed. But the shot is accurate because of the lucky second gust, with a distinctive luck that repels competence. (Sosa, 2021, p. 18)

However, this is more than a simple example among many others. Sosa offers us a general model for the normative evaluation of human performances, a model that becomes even clearer when the telic horizon of performance is also alethic – a normative model for evaluating human performances aimed at goals and, in particular, an evaluation of the action of the epistemic agent. Triple AAA, therefore, is a normative model, an almost grammatical criterion on identifying a result with the agent's merit, a measure of attribution of responsibility and the meaning of the act itself.

7. Telic normativity is inherent to actions, such as attempts that characterize human performances.⁵ It is telic because they aim at ends and often normative because we say they are better if successful and achieve their objective. It is also better for attempts to manifest competence and achieve success through competence, not by chance.

That is why we prefer persuasion to the use of force, an excellent diagnosis to mere guessing, the expert's advice to the charlatan's opinion. Also, we attribute merit to regular athletic performances rather than casual successes. After all, as Sosa reminds us, "to reach Larissa through ignorant luck is not to flourish." (Sosa, 2015, p. 142)

In light of a theory of telic normativity, the model accounts for performances in general and, in particular (as in the case of competing air gusts), of Gettieresque situations. The approach is, therefore, entirely normative, with prescriptions such as

1. It is better to be competent than inept;

5 "Telic normativity is inherent in action, in attempts. An attempt is "better" (as an attempt), a better attempt, if it succeeds than if it fails; better if competent than if incompetent; and better if its success is apt, through competence rather than luck. This all seems trivial for instrumental, end-relative normativity." (Sosa, 2023b, p. 189)

2. It is also better to be successful than not to be successful;
3. Even better is to succeed based on our merit, our competence, and how appropriate our action is.

A telic and epistemic gesture following this model can then be defined as:⁶

- a) **not being knowledge (or not being an act creditable to the agent), if:**
 - 1) $\sim$ Accurate, $\sim$ Adroit and $\sim$ Apt (like a blind shot that fails);
 - 2) Accurate but $\sim$ Adroit and $\sim$ Apt (like a blind shot that hits the target);
 - 3) Adroit but $\sim$ Accurate and $\sim$ Apt (like a shot made competently but which deviates from the target);
 - 4) $\sim$ [Adroit and $\sim$ Accurate, but Apt] (The negation of the conjunction indicates that it is a grammatical limit of the model (to use a Wittgensteinian expression), as it indicates something that, after all, can never occur. For it to be apt, it needs to be accurate and adroit as it is accurate because it is adroit.);
 - 5) $\sim$ [Accurate and $\sim$ Adroit, but Apt] (Another grammatical limit of the description);
 - 6) Accurate and Adroit, but $\sim$ Apt (like a shot that reaches the target, but due to the effect of double bursts that compensate each other, this being the truly Gettieresque situation of a starting competence that is not maintained upon arrival, as the accuracy does not express competence, and it is not open to debate whether the presence of luck, in any case, suppresses telic credit);

A telic gesture can be defined as:

- b) **being knowledge (or an act creditable to the agent) at various levels or gradations if:**
 - 1) Accurate, Adroit, and Apt (thus, basic level of knowledge, animal knowledge, either because it does not involve reflection or because it is unsafe);

6 With minor differences (perhaps improvements), we recover below our interpretation of the telic model as a normative evaluation criterion, as presented before in our text "A Gnoseologiasegundo Ernest Sosa" (Revista *Trans/Form/Ação*, v. 44, p. 63–96, 2021), from which we also took some of the following considerations about modal conditions for knowledge.

- 2) Accurate, Adroit, and completely Apt (that is, reflective knowledge, which can be safe knowledge at a higher level, with safety no longer being a necessary condition for knowledge but rather a distinctive note of the higher level of knowledge, a security knowledge).

We then have a triple standard for evaluating performances (AAA), to which, by the way, is added a triple component of evaluating adroitness, as the reliability of the competence depends on a more internal ability (skill, preserved even when we sleep), but which can be compromised by our effective state (shape) and by the concrete conditions in which the performance develops (situation).

Such a telic evaluation model also has practical applications. They are complex and far from automatic. However, the model serves as a nuclear formula for what can be charged to all human activity and what can be used for its evaluation, with the specifications and additions that prove necessary – sometimes, case by case. Such a model offers a broad perspective for evaluating performances, opening up an arc of questions that crosses every purpose-driven human endeavor. Therefore, Sosa states: “Our aim is a theory of knowledge that will fit a unified and broad-scope epistemology.” (Sosa, 2017, p. 210)

In this case, pragmatic aspects previously subordinated to epistemic aspects acquire unusual relevance to fulfill this broader purpose. It is no coincidence that Sosa's work, in its current progress, benefits from a more intense dialogue like Wittgenstein's *On Certainty*. As we will show in another paper, the analysis of human performances makes Sosa's reflection, together with Wittgenstein's, move to the concrete ground of actions involving language games amidst life forms. Then, it will come to the fore notions like “background conditions,” “hinge propositions,” and “bracketed domains.”

8. As presented in *Epistemic Explanations*, the normative model argues the merit or responsibility of the epistemic agent and the agent in general. Reliability indicates the probability of an attempt's success, given the model's components. That would be, let us say, a “safety condition” so that an attempt results in a hit, being harmless to luck. Let us quickly recover the reasons for introducing the modal notion of ‘safety’ for knowledge evaluation.

One way of analyzing knowledge was to propose a modal addition by which belief would acquire an uncontroversial characteristic note that would elevate it to the rank of knowledge. For this, for example, the be-

lief should be sensitive so that, if it were not true, we would not have it ($\sim p \rightarrow \sim Bp$) – if not p , then we would not believe that p . (Nozick, 1981)

Now, the condition of sensitivity is hardly sustainable and cannot resist skeptical threats. If there is *sensitivity*, if the facts were different, our senses (our gauges) would show something different. If what happens, if reality were different, the appearance would be different, and, therefore, our beliefs would be different.

However, the epistemic agent would be powerless in the face of the skeptic's claim that our experience can be illusory and misleading. In this case, our instruments would not indicate any change. They do not in themselves indicate, in the skeptical scenario, that they are disconnected, and we do not, after all, have access to the fact of access itself, as we have no connection with the connection – in short, we are not sensitive to the fact presupposed by the condition of sensitivity.

To avoid such an undesirable upshot, Sosa replaced sensitivity with a distinct epistemic modality, safety ($Bp \rightarrow p$). As sensitivity and safety are not equivalent, it seemed advantageous that novelty⁷ – namely, a belief capable of postulating the knowledge condition would be one such that we would have it only if it were true. Later, Sosa admitted that a belief can yet constitute knowledge, even if unsafe because it is overdetermined, but this applies only to the circumstance of so-called “animal knowledge.”

For instance, imagine that a baseball player's performance would be at risk because of the imminence of a power outage during a game. Can this risk deny the player's credit for hitting the ball? Can mere danger similarly suppress our epistemic capacity and our telic credit? Not at all. This confusion between knowledge and the presupposition of knowledge would lead us to suppress the possibility of human knowledge, its consequence being even more absurd, namely, the pure and straightforward assumption of no possible knowledge.

If we are lucky enough to escape this most devastating danger, this luck does not compromise our performance; in this case, we are not being reckless. For instance, the imminent risk of an electrical blackout, which was not confirmed due to extreme luck, does not suspend the credit of the victorious athlete hitting a ball.

As a result, Sosa has recently insisted that *not all luck corrupts a performance* – which, moreover, modifies the entire characteristic of the fa-

7 As Sosa explains in one of his most quoted texts, “[a] belief is sensitive iff had it been false, S would not have held it, whereas a belief is *safe* iff S would not have held it without it being true. For short: S's belief $B(p)$ is sensitive iff $\sim p \rightarrow \sim B(p)$, whereas S's belief is safe iff $B(p) \rightarrow p$. These are not equivalent, since subjunctive conditionals do not contrapose.” (Sosa, 1999, p. 146)

mous Gettier problem. So, we can see that the possible abandonment of safety as a necessary condition for knowledge affects only the b1 clause (related to animal knowledge), being still a normative condition for full knowledge (i.e., our b2 condition of an epistemic gesture Accurate, Adroit, and completely Apt).

9. Let us summarily present some final considerations, either derivative of what we presented or signaling future steps we could not present in this paper.
 - (1) Normativity is at the basis of epistemic justification, which, as such, can both explain what it is to know and *ipso facto* justify the relevance of explanatory undertakings.
 - (2) Normativity touches on modal issues essential to judging reflexively in a chain of knowledge classification. The hierarchy of knowledge does not have strict philosophical interest; nevertheless, being normative is far different from a mere taxonomy.
 - (3) The reflection ends by asking for more answers about the agent's perspective to identify whether justified, for example, in suspending an action – and suspending a judgment is an exercise that can never be reduced to a mere causal impulse.

Thus, Sosa's current reflection, which seeks to analyze the components of a foundation in primary conditions for knowledge, reiterates the normative features of his all-encompassing perspective while taking a new step in his epistemology of virtues – now extended, improved, and ready to become a Dawning Light Epistemology, which will soon entirely deserve our attention.

A highly respected scholar, Tim Crane, announced to the Four Winds in a recent tweet: "I'm rather surprised to say this, but I think Ernest Sosa may have solved epistemology." If Crane is correct (and we hypothesize that he is) and Sosa's proposal has the intended scope, the limits of possible experience will have been redefined from an epistemic point of view as if Sosa carried out a kind of new Cartesian revolution.

To carry out such a "revolution," Sosa does not disregard the fruits generated by the Gettiresque race in search of an analysis of sufficient and necessary conditions for what constitutes knowledge. However, by presenting a normative model, his theory goes beyond a definition capable of resisting the invention of any counterexample. On the contrary, it benefits from the very counterexamples that challenge it since it leads us to explanations of how what we call knowledge occurs. In its complexity, the

model situates different cases within its normative frameworks without imprisoning the different situations in the same straitjacket. The model allows us, on the contrary, a more flexible taxonomy through which the complexity of phenomena that involve knowledge can find its place within the framework of a well-established normativity.

Knowledge can be technically defined as an apt belief. In that case, this does not transform the work of analysis into the mere search for biconditional statements but instead allows the explanation of diverse manifestations and, therefore, of distinct classifications, which even incorporate different modal gradations. Thus, being clear about what knowledge is not, we are also better informed about what, for instance, animal knowledge, reflective knowledge, and full reflective knowledge can be.

In any case, true or false, Crane's statement is powerful. It is equivalent to claiming the proof of a theorem that has resisted for centuries or that a consolidated theory has been refuted. However, in this case, the statement is more comprehensive. It challenges all of us (or, at least, the science community) to answer what knowledge ultimately is and to draw from it the consequences for the valuation of human activity in general. Crane's claim is not self-proving; it does not follow from his academic authority, but it certainly deserves some attention, as we have tried to grant here.

References

- Gettier, E. (1963). Is Justified True Belief Knowledge?. *Analysis*23: 121–123.
- King, P. (2006). Anselm. In D. Borchert (ed.), *Encyclopedia of Philosophy*. Farmington Hills: Thomson Gale, vol. 1.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Plato. (2014). *Theaetetus*. Translated by John McDowell. Oxford: Oxford University Press.
- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (ed.) (1994). *Knowledge and Justification*. London: Dartmouth Publishing Co.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives*, 13.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Sosa, E. (2017). *Epistemology*. Princeton: Princeton University Press.
- Sosa, E. (2021) *Epistemic Explanations: A Theory of Telic Normativity and What It Explains*. Oxford: Oxford University Press.

- Sosa, E. (2023a). *Philosophical Methodology*. unpublished manuscript.
- Sosa, E. (2023b). Default Assumptions and Pure Thought. In L. R. G. Oliveira, (ed.) *Externalism about Knowledge*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Wittgenstein, L. (1975). *On Certainty*. Oxford: Blackwell.