



TOWARDS AN AUTHENTIC EPISTEMIC WORLD

*A Possible Worlds Solution to
Cognitive Dissonance and
the Gettier Problem*

CHRISTOPHER JAMES STEVENS

TOWARDS AN AUTHENTIC EPISTEMIC WORLD

*A Possible Worlds Solution to
Cognitive Dissonance and
the Gettier Problem*

CHRISTOPHER JAMES STEVENS



BORDEAUX
SELF-PUBLISHED

2017

*For Marion Vorms and Olivier Renaut, without
whose patient pedagogy and invaluable input my
philosophical world would be immeasurably poorer.*

Preface	IV
Introduction	7
1 Knowledge as an intellectual capacity	12
1.1 Virtue Epistemology	12
1.2 Knowledge that almost wasn't: when we know, but only just	20
1.3 Knowledge as dynamic manifestation of intellectual virtues	28
2 Resolution of Cognitive Dissonance	35
2.1 Cognitive Dissonance – the spontaneous relegation of JTBS	36
2.2 Coherentism or “bad faith”?	46
Conclusion	53
Bibliography	55

Preface

IT WAS IN 2012, while researching the present text, a sort of double-edged philosophical experiment in a new way of conceiving knowledge and a new way of doing philosophy, that my vision of conceptual systems first began to expand to global proportions. Less and less was I seeing them only as what their creators had no doubt intended them to be, *viz.* accurate descriptions of some aspect, if not the entirety, of our world. Instead, *beyond* those modest intentions, each individual conceptual system I meditated upon began to take on the dimensions of a distinct, self-contained, self-coherent, virtual or possible world all of its own: not a description whose trajectory began in Nature and ended in Mind, but an aesthetic creation drawn upon the formless whole of existence and projected back through the mind into a universe of its own, woven from its own fabric, suspended in the relationship space where the empirical 'sensible' world and the faculties of perception and understanding exchange and interface. Through this new mind's eye, the empirical world remained that infinite Heraclitean whole of constant flux, resistant to any form of static description from within, whereas all the possible, virtual worlds extruded from every conceptual system I considered, although still malleable and modifiable, could and would be just as fixed and immutable as their creator-inhabitants needed and willed them to be for any given expanse of space-time.

In such a vision, it was clear that the end of and means to *learning* a conceptual system would also need to shift. Assimilation of a conceptual system would now be complete only at the instant when the possible world this system gave rise to took full form between Mind and Nature, allowing the pupil to apprehend, sense, and spontaneously interact with the empirical world using this newly learnt possible world as a freshly acquired lens. What's more, just such an instant was something I had already experienced myself during my undergraduate years in Paris. After countless hours, stretched over two or three weeks, reading, re-reading, meditating over chapters from Kant's *Critique of Pure Reason*, one day, as a whole, it finally clicked into place and, quite literally, I suddenly and unexpectedly found myself interacting with a world whose every perceptible aspect had changed in both its sensorial character and conceptual implication. Most sublime of all, the experience unfolded as though necessarily accompanied by an intensely ecstatic physiological sensation. In later research I would discover many similar such vision-shifting testimonies from a whole range of thinkers

of various disciplines, almost every one of them reporting the same feeling of conceptual ecstasy.

Kant's philosophy, then, constitutes one such possible world, one which, from that fateful moment on, I could inhabit or step out of at will. In the first part of *Towards an Authentic Epistemic World*, we encounter two other such possible worlds: the world of propositional logic and Plato's world of Ideal Forms. In the second part, we are confronted with the putative existence of many, many more. Possible but virtual worlds, worlds which exist in reality only insofar as they serve as projected, focusing lenses for interacting with the infinite density of the empirical world. This, in any case, is how I conceive of them in the possible world of *my* philosophy, on which point I must happily concede the ouroboric nature of everything just such a philosophy implies.

In the years since that first moment of conceptual ecstasy, when I finally managed to see clearly through the lens of Kant's possible world, I have shifted my research focus to within the domain of the formal sciences and the conceptual systems which they comprise: Newton's possible physical world, Einstein's possible physical world, Darwin's possible biological world, Gould's possible biological world, etc. My aim is to study a particular kind of pluralism, to investigate the effects, both on our perception of Nature and on our brains and minds, arising from the superposition of two or more possible world conceptual systems in the relationship space between the empirical world and the organic functions of a given individual's understanding. One might call it a cognitive pluralism.

Towards an Authentic Epistemic World was neither knowingly nor intentionally begun as a 'pluralist' work. Originally, it was intended to be a philosophical project exploring the extent to which a boundary-free epistemic investigation could be achieved, one that would traverse borders of both discipline and denomination: examining questions raised by modern analytical philosophy through the lenses of both contemporary continental and ancient Greek and Roman philosophies, before obliging all of them to confront contemporary experimental psychology, endeavouring ultimately to steer the net result towards a positive, pragmatic, and eminently *applicable* normative proposition whose claim was to be as relevant to everyday life as it would be to any other level of epistemic investigation, be it personal, philosophical, or scientific. By the time the work was completed, however, it was clear to me that, in a very real way, I had cast a veritable constellation of variously angled possible conceptual world lenses, all mutually

PREFACE

refracting and re-focusing the images produced by their neighbouring worlds. However, in what, as I write, is now being labelled the *post-truth* world (although the text you are about to read will raise the contention that it would often be more accurately called the *post-impression-of-truth* world!), perhaps the most important figure manifest in *Towards An Authentic Epistemic World*, whether this be as the sports coach, the general, the psychologist, the philosopher, or indeed the author, is that of the educator: the voice occupied with reminding us, as much with every ‘success’ as every ‘failure’, that our ultimate endpoint should not be to win the race towards some ‘truth’, a goal forever liable to deceive us – reminding us, then, that the knowledge we have is less important than the thinking that brought us there. Indeed, our ultimate endpoint must always be the ceaseless sharpening of our cognitive and epistemic capacities: as humans intent on creating a richer global epistemic community, we must be forever learning, and teaching, how to *think better...* a lesson I believe the modern, academic philosopher audience the arguments to follow were originally aimed at needs to learn perhaps more than, and certainly *before*, any other group.

CHRISTOPHER JAMES STEVENS
Bordeaux, November 2016.

Introduction

ALTHOUGH the main body of this work involves references to specific perspectives within philosophy, its broader interest resides in the universal importance of its two central questions. Firstly, what is knowledge? And, once a satisfactory answer has been given to this first question, how should one best conceive of and acquire such knowledge? Or, to put it plainly, how should one best *think*? At the core of this exploration is a very short article by Edmund Gettier in which he demonstrates that knowledge can not just be the one thing many philosophers had thus far claimed it was: *justified true belief*. Initial responses to Gettier's now (in)famous article (Goldman, 1967, Lehrer & Paxson, 1969, Klein, 1971, to mention but a tiny selection) based themselves on a face-value reading, in the sense that these thinkers presumed Gettier's formulation of the epistemic problem encompassed it entirely. For one clear reason then, my justification for suggesting that they should in fact have broadened the scope of their responses is the observation that the core epistemic problem *does* go beyond Gettier's formulation of it. Indeed, the latter presents only one facet of what is a much deeper and genuinely metaphysical conundrum.

At first glance, Gettier's 'Is Justified True Belief Knowledge?' (1963) seems to show only that, just because any unit of knowledge must necessarily be a *justified true belief* ('JTB' from now on), this does not entail that JTB is a *sufficient* definition of knowledge (*i.e.* does not mean that JTB = knowledge), as, Gettier tells us, had been previously upheld (he cites Chisholm, 1957, Ayer, 1956, and Plato's *Theaetetus* as examples of this position). Indeed, it seems that Gettier presents nothing more than just this problem of a *necessary* but *insufficient definition*, whose solution, therefore, must surely be to simply modify and improve the JTB determination. The meat and bones of his concise article consists in the presentation of two case types wherein granting the status of 'knowledge' to certain specific JTB propositions, *P*, would be *logically* legitimate despite the fact that any human mind would naturally and intuitively conclude that the epistemic agent in question *does not know* that *P*. In the interests of reliability and transparency, let me quote the first of these examples in full:

*Suppose that Smith and Jones have applied for a certain job.
And suppose that Smith has strong evidence for the following
conjunctive proposition:*

INTRODUCTION

(d) Jones is the man who will get the job, and Jones has ten coins in his pocket.

Smith's evidence for (d) might be that the president of the company assured him that Jones would in the end be selected, and that he, Smith, had counted the coins in Jones's pocket ten minutes ago. Proposition (d) entails:

(e) The man who will get the job has ten coins in his pocket.

Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true.

But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. And, also, unknown to Smith, he himself has ten coins in his pocket. Proposition (e) is then true, though proposition (d), from which Smith inferred (e), is false. In our example, then, all of the following are true: (i) (e) is true, (ii) Smith believes that (e) is true, and (iii) Smith is justified in believing that (e) is true. But it is equally clear that Smith does not know that (e) is true; for (e) is true in virtue of the number of coins in Smith's pocket, while Smith does not know how many coins are in Smith's pocket, and bases his belief in (e) on a count of the coins in Jones's pocket, whom he falsely believes to be the man who will get the job.¹

To begin, let us focus on two elements in the actual wording of the article:

(1) A certain linguistic ambiguity whereby, on the one hand (notably in the title itself), Gettier employs the substantive form *knowledge* whereas, in the text itself, he employs the verb *to know*, without ever differentiating the act from its product; (2) Gettier's explicit use of the rules of propositional logic in the creation of his two eponymous 'cases'.

With respect to the first element, I will maintain that the philosophical implications of the intrinsic distinction and relationship between the verb *to know* (action) and the substantive *knowledge* (product) should have been explored as a priority from the time of the earliest responses. I advance that, by focusing on instances of *knowledge* or *belief*, we describe an insufficiency in the definition of a substantive, that is, some *thing*,

1. Gettier (1963), pp. 121-123.

whereas by concentrating on the mental capacities of *knowing* and *believing*, what reveals itself is a veritable metaphysical stumbling block whose implications go well beyond the definition of any single word or entity.²

When the problem is tackled as an issue with the status of *beliefs*, what is involved is only an examination of epistemic states. This is the approach we find in Goldman (1967) when he advances that what the JTB definition is missing is a demand for “appropriate causes” which would guarantee the relevance of the justification, as well as in Lehrer & Paxson (1969) and Klein (1971) when they propose, each in their own manner, a reinforced definition, the so-called GUJTB (*genuinely undefeated justified true belief*) definition.³ All such responses come under the umbrella of ‘JTB+1’ solutions to the Gettier problem. In other words, any solution that proposes a fourth condition to prop up JTB in order for it to become both a necessary *and* sufficient definition of knowledge. However, a wholly different approach is also possible, one which does not begin by examining an epistemic *state* but instead focuses on the epistemic *capacities* of the agent. This approach differs from JTB+1 solutions by moving aside static products (the epistemic states, *viz.* beliefs) in order to properly examine the machine that produces them. Indeed, just like a factory owner who, upon noticing a number of faults in his products, must seek the cause for these faults in the production line, in either the machines or his workers’ actions, so too this approach does *not* take on the task of fixing each defective instance of knowledge produced or, worse still, of redefining what the product is supposed to be to match the now faulty items. Instead, like the conscientious factory owner, this approach concentrates on the upkeep and correct functioning of the production line itself: the capacity for *knowing*. It has the advantage of focusing primarily on the epistemic agent *qua* agent, in the richest sense of the term, including all social and emotional motivations, cognitive capacities, and also whatever personal objectives are put to work by the individual in their quest to grasp hold of truth in any given, dynamic

2. The idea here is not to put words into Gettier’s mouth. Indeed, we can’t know if, for him, his cases constituted a simple situation of insufficient definition or a compelling reason to reconsider the metaphysical status of all knowledge. Whatever the case may in fact be, it is from this latter angle, we will see, that virtue epistemology tackles the article, using Gettier cases as one proof among many that indicate the need to embrace a new theory of knowledge.

3. Let it be noted in passing that, in the 1990s, Goldman eventually adopted the virtue epistemology approach (see below).

environment. The most well known contemporary representative of this approach goes by the label of *virtue epistemology*, but other schools also exist and the divergences between some of them will be explored over the course of this text.

The first section of this investigation lays out two quite distinct visions linked by just such a definition of knowledge: knowledge as a *capacity*. We will see that the ongoing debate internal to this conception of knowledge defines two opposing theoretical engagements; one which wishes to save and retain 'gettierized' knowledge⁴, the other which feels authorised and obliged to reject it. This debate seems inevitably to spill over into a more normative question concerning the function of knowledge in the very life of the epistemic agent: with respect to the epistemic agent and her relationship to her environment, should we demand a stricter definition of knowledge that would banish 'gettierized' knowledge from the accepted set? Or is a looser, more all-encompassing vision that would save this type of contested knowledge more desirable? To properly understand the implications of this question, I will borrow a conceptual analytical tool from the Stoic philosophical tradition of Antiquity: the fundamental distinction drawn between a 'goal' (σκοπός, *skopos*) and an 'end' (τέλος, *telos*).

The conclusion to the first section opens the way to a deep philosophical analysis of the second element mentioned above; Gettier's explicit use of the rules of propositional logic in the construction of his epistemic scenarios. Through comparison to the Platonic world of Ideal Forms, the world of propositional logic will be presented as one other kind of *possible world*^{*}, one Gettier draws upon in order to create his 'gettierized' knowledge, just as Plato rooted his conceptions of truth and justice, etc., in the ideal forms of these which properly existed *only* in the world of Ideal Forms rather than in our immediate, empirical, 'sensible' world. Bridging the first section to the second, we will then see how this shared Gettierian/Platonic method of exploiting *virtual*

4. The notion of 'gettierized' knowledge will become clearer as we progress. In a nutshell, it refers to any instance of justified true belief produced in a Gettier style situation; a logically legitimate JTB which, based on certain circumstances of its production, one may intuitively feel is not, in fact, an instance of knowledge. "The man who will get the job has ten coins in his pocket" is one such example of 'gettierized' knowledge.

* The original French version of this article was written before I became aware of Nelson Goodman's possible world philosophy as exposed in his excellent *Ways of Worldmaking* (1978). Given the nature of his thesis, our two visions end up being entirely complementary, and I recommend his book as a perfect text for further exploration of the perspectives set out here.

worlds, including all the epistemic problems this entails, can in fact be seen as a recourse that is universal to human reasoning.

Hence, in the second section, we will analyse the *effective* function of knowledge within the real (empirically testable) cognitive life of the agent. This analysis contains two complementary perspectives, one psychological, the other philosophical, and will once again lead us towards a discussion of the actual, real-world relationship between belief and knowledge, this time, however, from an experimental and theoretical point of view rather than an analytical one. We will see that the cognitive agent, in the course of her everyday life, blurs any possible distinction between belief and knowledge and, in order to maintain maximum cognitive *coherency* (Harman's vocabulary, 1986) or *consonance* (Festinger's vocabulary, 1957, 1959), she is liable to quite blatant replacement of genuine JTBs for conjured up beliefs of very weak justification and highly dubious veracity.

Resulting from this discussion, and serving as a conclusion to the investigation, is a demonstration of the hypothesis that our *knowledge* is nothing more than a nominal label we give to a certain type of *belief*. This knowledge-belief derives its seemingly permanent, absolute epistemic value from the rigid coherence it draws from cognitively generated possible or virtual worlds (in the sense that Propositional Logic and Pure Mathematics constitute virtual worlds), worlds which can only ever partially correspond to the empirical world we actually inhabit and live in. We will see that the resolution of cognitive dissonance often involves a seemingly unconscious and yet abusive use of this capacity for creating other, merely *possible* worlds and that this behaviour very closely resembles what Sartre refers to as the conduct of "bad faith." Based on this proximity, I propose that a form of Sartrean authenticity (*i.e.* Sartre's solution to "bad faith") should be added to the nascent virtue epistemology outlined in the first section in order that the importance of avoiding abuse of the capacity for virtual world creation – an abuse underpinning both the Gettier problem and the quasi-universal "bad faith" resolution of cognitive dissonance – be openly acknowledged.

Knowledge as an intellectual capacity

IN THIS SECTION, I will lay out the key features of virtue epistemology as it has grown up around both Ernest Sosa's seminal article 'The Raft and the Pyramid' (1980) and certain ideas it has adopted from Roderick Chisholm's *Theory of Knowledge* (1966). Virtue epistemology was certainly not conceived of with the aim of solving the Gettier problem, but from its very foundations arises a change in perspective that affords a whole new way of perceiving it. This new perspective will be teased out as we proceed through this section. The second objective will be to present a distinction Angélique Thébert (2012) has drawn between the theoretical concepts of intellectual *capacity* and intellectual *virtue* (as this is understood in virtue epistemology). This critical distinction and its accompanying 'relaxed' vision of what should be accepted as knowledge will then be critically assessed in turn. Following a conceptual analysis (borrowed from the analogy of the archer found in Cicero's writings on Stoic philosophy) of Thébert's thesis, I will present a new, more dynamic and nuanced conception of knowledge inspired by the work of John Turri (2012). This conception seems to get past Thébert's objections while also demanding a modification to virtue epistemology's overall approach. Finally, a passing reference Gettier makes to Plato will be closely scrutinised, perhaps for the first time, revealing much more than previously suspected about the metaphysical origins and impact of Gettier's famous problem.

1.1 Virtue Epistemology

Let us begin with a definition virtue epistemology gives us for what a virtue is: "A virtue, in one important sense, is an ability. An ability, in turn, is a stable disposition to achieve certain results under certain conditions. Further, when we say that a subject *S* has an ability to achieve certain results, we imply that it is no accident that *S* achieves those results."⁵ The most essential dimension of this definition is that, in making an ability or capacity of virtue, it thereby binds it to an essential characteristic of any capacity: *reliability*. To illustrate this, John Greco (1993) evokes the image of a sportsperson's capacity to execute certain

5. Greco (1993).

actions related to their specific sport: Rafael Nadal has the capacity to execute a serve, Zinédine Zidane to accurately pass a football, and so on and so forth. Thanks to a specific set of properties – both innate and acquired, internally and externally derived – all accomplished sportspeople have a tendency to succeed in these actions in normal conditions. It is to be noted, however, that this capacity does not guarantee infallible success: even under normal conditions, a mistake may occur. The important implication of this for virtue epistemology, as a *normative* discipline, is that the manifestations of a virtue can forever be improved upon. The ‘rules’ of ‘good thinking’ that virtue epistemology is liable to propose will therefore be like archetypes, cognitive ideals to be aimed for, enabling us to affirm that this is a normative school, insofar as it identifies a specific inherent property of intellectual virtues: they are *always* perfectible. Another, more succinct definition, simply says that epistemic virtues are “dispositions that make the subject good at detecting and endorsing the truth, so that she usually gets it right when she exercises those dispositions.”⁶ What does such a conception imply for a theory of knowledge?

First, let us recall that virtue epistemology was born out of the application of certain theoretical ideas from virtue *ethics* to the study of belief formation and, more specifically, out of Ernest Sosa’s attempt to overcome the conflict between the coherentist and foundationalist⁷ responses to the question of knowledge justification. The influence from theories in virtue ethics is most prominent in the following two facets of virtue epistemology. Firstly, the relation of supervenience sought between the evaluative justification quality of a belief and the non-evaluative properties (such as logical implication, probability, etc.) upon which this justification supervenes. In explanation, Sosa (1980) affirms that, “It is a goal of ethics to explain how the ethical rightness of an action supervenes on what is not ethically evaluative or normative. Similarly, it is a goal of epistemology to explain how the epistemic

6. Turri & Sosa (2013).

7. The article ‘The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge’ (1980) constitutes, in fact, a long debate between these two positions without, however, proclaiming either one of them as victor. What is needed to surpass this no-win debate, according to Sosa, is a new teleological approach whose aim would be the future maximal epistemic gain of a given epistemic community. He calls this approach, in section 11 of his article, *reliabilism*: “In epistemology, there is reason to think that the most useful and illuminating notion of intellectual virtue will prove broader than our tradition would suggest and must give due weight not only to the subject and his intrinsic nature but also to his environment and to his epistemic community” (p. 23).

justification of a belief supervenes on what is not epistemically evaluative or normative.”⁸ Such a goal contains echoes of Alvin Goldman (1979), who also wanted to find a theory of justification that would specify what a justified belief was but in exclusively *non*-epistemic terms. To explain this somewhat counter-intuitive idea, Goldman and Sosa both draw an analogy between the epistemological notion of justification and the ethical notion of ‘good action’. To take just one example of this, utilitarian theory deconstructs the notion of ‘good action’ by defining it as that action which produces, or would produce, at least as much total happiness as any other action available to the agent. So, in theory at least, there is no qualitative evaluation but, instead, only quantitative calculation. The question to be posed thus shifts from “Is this action good?” to “How much happiness does this action produce?” Hence, virtue epistemology relates to ethics precisely by its own parallel goal: to evacuate epistemic terms (like *justified*, *warranted*, *has reason to believe*, *knows that*, *sees that*, *shows that*, etc.) and thereby shift from the question “Is this belief justified?” to a question of the type, “How many perceptible proofs are there to support this belief?” To say of someone “She *believes that P*” is epistemically non-evaluative for the reason that no *value* is attributed to the belief. There is only the observation that some belief is in existence. In contrast, to say of someone “She *has reason to believe that P*” is epistemically evaluative for the opposite reason: here we have a clear attribution of truth value. So virtue ethics and virtue epistemology have this (at least) in common: they both aim to ground the objects of their respective disciplines (goodness, morality, knowledge, etc.) in the *non-evaluative conditions* of these objects. Sosa (1985), summing up this non-evaluative aspect, says: “if a belief is cognitively justified, it is so presumably in virtue of nonevaluative properties, perhaps having a certain source in perception, introspection, memory, or inference, or some combination of these. And any other belief just like it in respect of all such properties could not fail to be equally well justified.”¹⁰ Furthermore, in considering the quality of an action, virtue ethics takes into account the agent’s *disposition* while the action in question is being carried out. For example, can we really speak of a ‘good action’ when, by pure chance, some happy event results from ill intention? In other words, from the action of a badly disposed agent, or one employing some virtue that is unfit for goodness? And likewise,

8. Sosa (1980), p. 16.

9. List abridged from Goldman (1979).

10. Sosa (1985).

if an agent is not properly disposed to gain knowledge of something, can we say that she nevertheless still acquires some by chance? No, say the virtue epistemologists, precisely because the inclination towards chance or luck is not an *improvable* capacity and therefore can have no place within a project to develop a normative discipline. The specific aim of such a project would be to improve the agent's cognitive virtues by discovering the norms of reasoning, something luck and chance can, for better or worse, play no role in.

To sum up: taking its cues from virtue ethics, virtue epistemology seeks the classification of beliefs both through their justification, calculated according to epistemically non-evaluative properties, and through consideration of the dispositional conditions under which the agent produced the beliefs in question. Why these specific stipulations? It is hoped that the first will provide a certain epistemic coherence for the classification of beliefs (recalling that "any other belief just like it [...] could not fail to be equally well justified"), while the second, relating to the agent's disposition, will ensure that these beliefs stem from fitting cognitive virtues. This is equivalent to stating that, in order to be classed as knowledge, a belief must, of course, meet certain objective and non-evaluative criteria of justification (of the type just mentioned), but must also, above all, be the product of a cognitive virtue, whose specific and reliable function is to search and quest for truth and which is successfully mobilised towards the accomplishment of one goal: the acquirement of truths (which also entails that the truth of a belief must absolutely result from cognitive virtues, an aspect whose pertinence will become clear below).

The starting point of any virtue epistemology, as we have just seen, is the assertion that the discipline is a normative one and that, therefore, a purely scientific or descriptive analysis of cognition cannot suffice: there will always be something that escapes such an analysis, especially when it comes to what *good thinking* is. Virtue epistemology, through and through, in its goals and in its discourse, wears its badge of normativity with pride. Indeed, its problematic could be formulated like this: Given the cognitive mechanisms of reasoning and everything else we know about humankind and its environment, how *should* we reason?

For the virtue epistemologist, a cognitive capacity that usually succeeds in producing knowledge (and not merely true beliefs) would then be the 'best' model, the most *virtuous* model and, therefore, the model that one should emulate when reasoning. Related to these

goals is the second point all virtue epistemologies have in common: they all revolve around agents and cognitive communities and take these as the measure in accordance with which all epistemic values and evaluations must be decided, as opposed to other approaches that claim to determine the epistemic value of beliefs by analysing only their propositional coherence and content correspondence to 'reality', just as one would determine the truth values of propositions in logic¹¹: "...virtue epistemologists agree that the ultimate source of epistemic normativity, and hence the central focus of epistemological inquiry, are cognitive agents and communities, along with the fundamental powers, traits and habits that constitute their intellect. This contrasts with the mainstream approach in later twentieth-century analytic philosophy, which focuses on individual beliefs and inferences, instead of individuals and their cognitive character."¹² In a word, for any virtue epistemology, the cognitive agent is the condition *par excellence* of all epistemic questioning and all solutions to the 'how to think' question must therefore begin and end with her. We should also note that virtue epistemology considers the cognitive agent to be the collection of her virtues *and* her vices, something which solidifies the above demand that a belief *P* and a belief *Q*, even if they be identical in terms of their propositional content, are not *epistemically* equal if *P* was produced by a rigorous cognitive process whereas *Q* was the result of pure chance, foolhardiness, or indeed the result of a Gettier-esque improbability. For virtue epistemologists, this type of intellectual vice must not be nourished.

Therefore, it is clear that virtue epistemology falls under the umbrella of those theories of knowledge which *reject* 'gettierized' knowledge and which, more generally, seek to narrow the borders of what should be counted as genuine knowledge. Thus, the process of knowledge production becomes a norm, an ideal to which we should aspire in exercising our cognitive capacities. This normative aspect of virtue epistemology inevitably opens a horizon representing its *responsibilist*

11. Note that this is precisely the nature of Gettier's own sleight of hand: sliding from the real world into the world of propositional logic – as evidenced by the vocabulary he employs ("conjunctive proposition," "implication," etc.) – in order to give a truth value of 1 (*i.e.* 'true') to his proposition (e): "The man who will get the job has ten coins in his pocket." He then attempts to return to the real world with this truth value still intact. Yet, there is a genuine and strong disagreement between these two worlds, one which we will explore in more detail further on.

12. Turri & Sosa (2013).

side. Indeed, if the criteria for earning the title of 'knowledge' are to be made stricter, one has the right to ask *why* this should be done. For many, the only full answer resides in the portion of social responsibility each and every agent carries with them as a member of a global epistemic community: we *should* demand this manner of reasoning in the production of knowledge because we *should* want to lead the community and each of its members towards truth. Otherwise, no communicated belief would be reliable. Hence, if reliabilist virtue epistemology is normative, then its responsibilist sibling is doubly so. Nevertheless, in the optic of a study targeting only Gettier cases, we can somewhat gloss over this global and social aspect – it is enough for us to simply presume that the cognitive agents we will be analysing are genuinely motivated towards finding truth rather than falsity. In practical terms, whether this motivation be of a purely individual nature or whether it be reinforced by some weightier social pressure, this will have little effect on our specific line of research here: in both cases the target is truth and virtue epistemology aims to tell us what we should do in order to attain it. Precisely, virtue epistemology's fundamental enterprise is to study the optimal functioning of reasoning in order that it may prescribe a normative epistemic ideal to the agent. Ultimately, it is this very intent that enables us to analyse Gettier cases through the lens of virtue epistemology and then to say whether the cognitive capacities employed by Gettier's agents fit or do not fit within that ideal. Or, in other words, to say whether these agents hold on to their JTBS *thanks to* their cognitive virtues or just *because of* some other (lesser) reason. So now, let us begin this analysis with another Gettier case¹³:

HUSBAND: Mary enters the house and looks into the living room. A familiar appearance greets her from her husband's chair. She thinks, "My husband is home," and then walks into the den. But Mary misidentified the man in the chair. It is not her husband, but his brother, whom she had no reason to think was even in the country. However, her husband was seated along the opposite wall of the living room, out of Mary's sight, dozing in a different chair.

13. This Gettier case comes from Zagzebski (1996) and is, epistemically speaking, identical to Chisholm's more famous Gettier case concerning a dog in a field which a farmer mistakes for a sheep (1966, ch. 6).

Firstly, let's start by identifying the element in this scenario which is common to all Gettier cases: the *double-luck* element. We have an instance of bad luck (Mary is mistaken in her identification of her husband) whose negative epistemic consequences are immediately cancelled out by an instance of good luck (her husband is nevertheless actually in the room). This double-luck property of Gettier cases, and the various ways in which it can be understood, will play a central role throughout the rest of this study. So why, according to virtue epistemology, does Mary *not* know that her husband is in the living room? Several answers, each closely related, are possible:

I. Mary arrived at a certain belief, which is true, thanks to her cognitive capacities (the inference she made was totally acceptable). She did not, however, arrive at the *truth* thanks to these virtues, for the truth of the matter was the result of chance, an instance of *good luck*.¹⁴

II. In order to *know* that her husband was in the room, the *credit* for this knowledge would have to be attributable to Mary. However, the credit for knowledge is only attributed when (1) the agent believes in the truth of a proposition thanks to her cognitive capacities, which must form a necessary and significant part of all the factors which led to the true belief, and when (2) no other factor can 'trump' the explanatory salience of these capacities. Yet, it is precisely the case here that the bizarre confluence of epistemic bad luck and good luck which led to Mary's JTB obviously trumps her cognitive capacities in terms of her conclusion turning out to be 'true'. Thus, she can not be accredited, thus she does not *know* that her husband is in the room.¹⁵

III. All evaluable performances, including beliefs, can be subjected to an AAA evaluation¹⁶ which measures how *accurate*, *adroit*, and *apt* they are. An accurate performance accomplishes its goal, an adroit performance displays competence, and an apt performance is accurate *because* adroit. Pertaining to beliefs, accuracy relates to truth, being adroit is equal to manifesting intellectual competence, and, finally, a belief is apt if it is true *because* intellectually competent. Using such an evaluation, we quickly see that the accurate *truth* of Mary's belief owes nothing to her intellectual competency, *i.e.* her belief was not *adroit* and therefore not at all *apt*. Rather, this truth was the result of pure chance and thus she does not *know* that her husband is in the room.

14. Zagzebski (1996).

15. Greco (2003).

16. The AAA method of evaluation and the content of this third response are from Sosa (1991, 2007, chs. 2, 4 and 5).

In reality, these three responses constitute three variants of one and the same vision: a belief is knowledge if and only if the truth of that belief comes to the agent thanks to the correct exercise of her cognitive/intellectual virtues. Indeed, the variations in the responses stem only from differing manners of defining the notion of 'correct exercise'.

Nevertheless, in order to more rigorously test this tripartite evaluation, let's sprinkle a little classical skepticism into the mix and imagine the following case, inspired by the 1975 film *The Stepford Wives*:

ROBOT: Mary, a paranoid woman who lives in constant fear of the worst, gets home – worried – and looks into the living room. A familiar appearance greets her from her husband's chair. She thinks, "My husband is still alive!" and then, somewhat relieved, walks into the den. But the man in the chair is not her husband, it is a robot in the perfect image of her husband, practically impossible to distinguish from the real flesh and bone man. Her husband, in fact, has been imprisoned in a dungeon for 10 years, alive, certainly, but far from the eye of his wife, who suspects nothing.

In **HUSBAND**, once we had given it the virtue epistemology evaluation treatment, it seemed intuitively correct to refuse the title of knowledge to the lucky truth Mary had arrived at. And yet, in **ROBOT**, a case which is specifically imagined to be epistemically identical to the former, when we apply the same evaluation techniques and requirements, we clearly begin to tread into the slippery terrain of demanding that instances of knowledge be absolutely *indubitable* beliefs. What would be the *correct exercise* available to Mary in order for her to arrive at the truth of this second situation? Which intellectual virtue should she have used in order to even suspect that truth may be absent, before even thinking about setting out on a quest to find it? Standard virtue epistemologies are left wanting before such questions. For, without access to some external and omniscient vision of things, by what means can we *know* that some true belief or another is *indubitably* true and, therefore, a genuine instance of knowledge? Virtue epistemology now seems to put us in a situation where knowledge *can* be arrived at, but only at the end of some infinite skeptical quest where everything is questioned by proper AAA use of our intellectual virtues. However, while classical

virtue epistemology approaches may seem thus far unable to answer this criticism, the horizon Gettier cases share with classical skepticism has not escaped the attention of everyone, and so this union deserves closer investigation before we see how contemporary virtue epistemology may itself respond to the challenge of a forever *dubitable* epistemic world.

1.2 Knowledge that almost wasn't: when we know, but only just

In the introduction, I evoked two sets of responses to the Gettier problem. We have just seen one version from the set whose aim is to deny the status of 'knowledge' to any 'gettierized' JTB. It is time now that we look to the second set, the smaller of the two, containing those responses proposing that *certain* 'gettierized' JTBs do in fact constitute legitimate instances of knowledge. An argument at the forefront of this conception can be found in the 2012 article 'La connaissance comme capacité. Le traitement des cas de Gettier' ('Knowledge as a capacity. Dealing with Gettier cases') where author Angélique Thébert (highly inspired by Stephen Hetherington, 1998, 2001, 2006, 2008¹⁷) develops the quite ingenious idea of taking certain arguments, usually employed to counter classical knowledge arguments from skepticism, and employing them to help solve the Gettier problem. The preface to her project nevertheless goes to pains to distinguish between Gettier's 1963 article and the usual skeptic stance. In it she states that, "the primary concern in Gettier situations is the *definition* of knowledge, not its *possibility*."^{18,19} Thébert affirms that, this distinction notwithstanding, there is nevertheless a certain family resemblance between the structure of the skeptic doubt argument²⁰ and the structure of the Gettier argument against the JTB definition of knowledge. So that everything is quite clear, this family resemblance between the skeptic vision and the vision

17. As this text was originally written in French, this explains why I chose a French representative of this point of view to analyse. However, my arguments, as indicated, could just as well counter Hetherington's position. Finally, Thébert does introduce some novel angles of her own, making her theses worthy of wider critical attention.

18. Thébert (2012), p. 1. The translation is mine, the italics the author's.

19. A reminder here that my own aim is precisely to contest any peremptory conclusion tending to presume that Gettier cases do not throw into doubt the possibility of 'knowledge', as we commonly understand it, in the empirical world (*i.e.* a subject's accurate belief relative to some objective truth). This is all the more pertinent to highlight given that Gettier never actually gave any indication as to what his own "primary concern" was.

20. In a nutshell, the basic Cartesian doubt argument: wherever epistemic conditions can be doubted, 'knowledge' is *ergo* impossible.

of those philosophers who, with Gettier, refuse to accept ‘gettierized’ JTBS as instances of knowledge resides in a shared demand for epistemic absolutism: either there is knowledge or there is not knowledge, no middle ground is accepted. Thébert, on the other hand, wishes precisely to “account for the gradualism of knowledge” and refuses to accept that knowledge is an “all-or-nothing” affair. But is it true to say that virtue epistemology also proceeds in this same absolutist fashion? Let us look at the arguments Thébert presents in support of her position.

The first step of her argumentation consists in differentiating two distinct epistemic structures – Gettier *situations* and Gettier *cases* –, the purpose of which is to bring some much needed clarity to a debate often dominated by confused overlappings. Thébert defines a Gettier *situation* as a situation in which an agent “believes that a proposition is true on the basis of reasons which are adequate from her point of view but which do not match the reasons for which her belief is in fact true.” Gettier *cases*, on the other hand, refer only to “[those] Gettier situations in which JTBS fail to be instances of knowledge.” Precisely, the objective Thébert gives herself in her article is to show that a Gettier *case* does not *necessarily* result from every Gettier *situation* and that, consequently, we can and should bestow the title of ‘knowledge’ upon certain JTBS which, until now, have been the baby in the bathwater virtue epistemology has been urging us to throw out. In other words, and it is an important point for her argument, Thébert accepts that *some* Gettier situations really are Gettier cases, that *some* ‘gettierized’ JTBS are *not* instances of knowledge. Notably, she agrees that HUSBAND constitutes a Gettier *case* because the truth of the situation is not related to any act of the agent’s. For Thébert, this is just one example of those Gettier situations she calls *unlucky*, that is, those where no knowledge obtains. Instead, what she proposes for our consideration are *lucky* situations such as the following (originally found in Russell, 1948):

CLOCK: Passing in front of a clock, I look at it and notice it reads eight o’clock. I conclude that the time is eight o’clock. However, without my knowing it to be the case, the clock does not actually work, its hands came to rest in the eight o’clock position several days before. Nevertheless, and by some lucky chance, it just so happens that it really is eight o’clock when I look at it.

The fact that the non-functional clock is displaying the correct time is mere coincidence. Had I passed by at any other time of the day, I would have been incorrect in concluding that it was eight o'clock. Thébert, however, proposes another way of considering things by focusing on the agent's (the 'me' in the above illustration) own epistemic situation (rather than on any external and omniscient view such as Gettier grants himself relative to Smith in the very first example we saw). I consider that I really do know that *P* (that it is eight o'clock) since I have no idea that the clock is not working. The situation in which I find myself is indistinguishable *for me* from the situation in which the clock works perfectly. I am not like Mary in **HUSBAND**, whose error was to mistake another man for her spouse. I have made no such mistake – I saw and read the time indicated on the clock properly. I know that *P* and even if I *almost* didn't know that *P* (had I passed by any earlier or later), the important thing is that I do actually know it. On this basis, Thébert rejects the approach taken by the virtue epistemologists because she finds it to be static, stemming from the fact that "according to them, knowledge is a true belief *produced by* a virtue or a capacity [...] whereas, in the alternative conception that [she proposes], the capacity defines knowledge, it has a *constitutive* role." And it is for this reason that virtue epistemology must exclude the fact that "I know that *P*" even though the common sense view, Thébert asserts, would be to accept this fact.

This approach borrows heavily from Stephen Hetherington's ideas (1998, 2001) about the role luck (whether this be *good* or *bad* luck) plays in the constitution of knowledge, something he refers to as "knowing luckily." The basic idea is that, just as figure skaters or tennis players sometimes execute their winning manoeuvres with the fortuitous aid of chance, "as if by a miracle," so our intellectual capacity has the right to claim the production of genuine knowledge even when it has been 'won' fortuitously (like in **CLOCK**), and this for the simple reason that we already evaluate our *other* capacities in just such a graduated way. Indeed, if Rafael Nadal scores a point during a tennis match by means of the exceptional execution of some foolhardy manoeuvre which his coach, in his infinite wisdom, has formally forbidden him to use, the umpire will not, however, deny him his right to this point. Game, set and match to common sense it would seem. However, before we arrive at any too-hasty conclusions, what Thébert intends when she says that knowledge is not produced *by* intellectual capacities but just *is* such a capacity is perhaps not yet sufficiently clear. So let us take a moment to examine the claim.

Taking the distinction Gilbert Ryle (1945) makes between propositional knowledge ('knowing that') and knowledge of the 'knowing how' variety in order to consider it in the light of Hume's thesis on the origin of inductive inferences, Thébert arrives at the following conclusions. Propositional knowledge such as "knowing that *P*" presupposes the possession of certain intellectual capacities, certain types of 'know-how' whose execution requires no understanding whatsoever of the operation's detailed workings. Hume, in his *An Enquiry Concerning Human Understanding*, a source Thébert quotes from, relates the development of these intellectual know-hows exclusively to habit: the habit of observing a uniformity in the same effects obtaining from the same phenomena, *i.e.* a level of inference even animals and children are capable of. This, of course, does not suffice to conclude that there exists no general principle of inference at work in the background, but it does show that "inductive know-how does not occur due to consideration of such a principle" (Thébert).

Coming back now to the situation in *CLOCK*, the author correctly remarks that virtue epistemology would not take 'my' belief that it is eight o'clock to be an instance of knowledge, for the reason that I did not arrive at that truth through only my cognitive virtues. Indeed, the improbable coincidence of the clock ceasing to function at an instant where its hands indicate the same exact time as I then happen to look at it some days later *trumps* the use I make of my cognitive virtues. But, says Thébert, wishing to deny such an instance of knowledge stems from two strong principles upon which virtue epistemology rests, principles which, in her opinion, are in fact *too* strong:

Principle of security: If one knows that *P*, then one could not have been easily mistaken in regard of some similar case.

Thébert explains the implications of this in the following way: "This principle highlights that, if knowledge does obtain, then it must be certain, it must be stable enough to be not easily overturned." From here, she then adds that this first principle leans on another, even stronger still:

Principle of infallibility: Knowing that *P* is knowing that all the possibilities of error associated with *P* are false.

And this is where we realise that we have been led, cunningly and adroitly, into the very heart of pure skepticism. Indeed, what the principle of infallibility seems to demand for genuine knowledge to obtain is nothing less than the omniscient capacity to exclude *every* eventuality whereby some error might obtain... including the possibility

of robo-husbands, of dogs perfectly disguised as sheep, and even of the bizarrely regular occurrence of clocks ceasing to work in such a way as their hands remain frozen at the time it happens to be at the instant I later walk by! Such a demand is so strong that, brought to its extreme logical conclusion, it would exclude the possibility of any certain knowledge whatsoever. On this basis, Thébert concludes that virtue epistemology is, perhaps unintentionally, inspired by Cartesian skepticism when it considers that JTBs of this kind are not instances of knowledge, claiming that they are not representative of a normal world: when I look at a clock, and before concluding that the hour it displays is the correct hour, would virtue epistemologists really have me stare at it long enough to check that the hands are actually moving, then double-check the time indicated on at least one other clock before, finally, scouting around the whole vicinity to make sure no evil demon is lurking about playing solipsistic tricks on me? For Thébert, the distinction that virtue epistemology makes here is as arbitrary as it is convenient. According to her, virtue epistemology wants to have its cake and eat it: virtue epistemologists wish for knowledge to imply a certain value and so, driven by that desire, they develop demanding criteria for the obtainment of the status of knowledge. Yet, at the same time, they also wish to affirm that we live in a *normal* world where the correct usage of intellectual virtues reinforces and improves them. “But where,” Thébert asks, “do we draw this demand for normality from?” Her opinion is that Gettier situations are far more common than one would tend to think, that we often find ourselves in situations where “the reasons for which we think our beliefs are true do not at all match the reasons for which our beliefs are in fact true.” Even if we admit this point (although Thébert does little to back it up), is it therefore necessary to consider these beliefs to be instances of knowledge?

Up to this point, the family resemblance identified by Thébert between philosophers who refuse ‘gettierized’ JTBs and Cartesian skeptics seems to be adequately justified. However, there is another such resemblance, which seems all the more striking and pertinent, but which she does not evoke. This resemblance links virtue epistemology’s approach to knowledge and good thinking to the distinction employed by the Stoics of Antiquity (sworn intellectual adversaries of the Skeptics, as it

happens) in order to separate the notions of goal (*skopos*) and end (*telos*).²¹ And yet, once this comparison has been made, it becomes difficult to read Thébert's article and not conclude that she has understood neither the veritable goal of virtue epistemology nor what it means when it describes itself as *normative*.

The classical analogy used to illustrate the *skopos/telos* distinction fits perfectly with some of the sporting analogies we have already seen above: it is the image of the archer in training. The archer's *telos* differs from his *skopos* insofar as the latter consists only in whether or not a given arrow's flight brings it to strike its target, whereas the former consists in the archer's very capacity to *aim well*. The 'good' archer must attach as little importance as possible to the *skopos* for the very reason that the fact of the arrow actually hitting or not hitting its target is susceptible to too many external interferences over which the archer has no control: a gust of wind could suddenly arise after he has released his shot and change the direction of the arrow, or else an obstacle, such as a bird, could suddenly swoop in and be struck by the arrow, preventing it from reaching its intended target, etc. Cicero summed up the distinction in this way, "the good archer is not he who hits the mark through luck, it is he who did everything to hit it that depends upon a good archer, even if he then miss it through the concurrence of events which are independent of him or which he could not have predicted."²² Likewise, the good tennis player will not sing his own praises over the shot he scored by sheer fluke, rather he will scrutinise himself with the 'omniscient' eye of his coach and get back to his training as soon as possible: the *end* he is aiming for is not simply to win but rather to be the best player that he can be. Certainly common sense, upon which Thébert ultimately leans her entire argument, will celebrate the player for every shot he scores, regardless of whether the means employed to attain them are perfect or not. But the Stoic outlook knowingly, wilfully distinguishes itself from common sense, aiming to raise itself above it. Similarly, in declaring itself a normative discipline, virtue epistemology aims for the creation of thinkers who are *good* in an identical sense, thinkers whose aim is not to be able to cry "Knowledge!" as often as possible, as if they were engaged in some kind of game between savants, but thinkers who measure themselves against only their own capacity and potential

21. To be clear on this point, I am not relating virtue epistemology to the Stoics' own unique theory of knowledge. I take only their *skopos/telos* distinction and give it a novel and, I believe, fruitful application.

22. *De Finibus*, III, 6, 22.

and so always seek improvement, raising themselves above the laurels handed out by common sense. Thébert never mentions this normative aspect of virtue epistemology. It is as though, both having observed that the world is largely governed by random chance, the Stoics conclude that everything must be done to try to isolate our capacities from its effects while Thébert concludes that we must make room for chance in our epistemic judgements.

At the end of this analysis, we can respond to Thébert that, in her criticism of virtue epistemology, she has missed her target not due to bad luck but because she was aiming badly. Playing the good, wise Stoic, we would say to her that the philosopher's role should be like that of the coach: not to congratulate the player because the ball happened to fall where it was supposed to, but to take this external and 'omniscient' point of view which *does* see everything and therefore knows what remains to be improved. Recognising knowledge to be a capacity should have been only a first step towards something richer: this capacity still needed to be separated into its distinct *skopos* and *telos* parts. It should now be clear that, for Thébert, the capacity to know remains a *skopos*-only question while for virtue epistemology it is much more a question of the *telos*, even if this specific terminology has never yet been employed by virtue epistemologists themselves. We can agree with Thébert that knowledge should not be seen simply as a product. But then neither should the *skopos*. It is a full process, a full movement: the arm pulling, the arrow flying, the target hit. This, however, is not the problem. Just because we acknowledge a constitutive role for intellectual capacity in the production of knowledge, this does not mean that we have chosen the *correct* role. When Thébert concludes her article by saying, "If knowledge does have a distinctive value then it is its gradualism, its magnitude (which is specific to the capacity [*i.e.* knowledge] but lacking in the belief) that enables it to have it," my Stoicism infused response must be to say that, in fact, for the *good thinker*, 'knowledge' can have no intrinsic value of its own since it is *only* the capacity to know which has any worth, and this for the precise reason that it is the only aspect of knowing which is not subject to external happenstance. Devoting oneself to raising the value of any given *skopos* can only result in diminishing one's *telos*, especially when one's commitment to that endeavour leads to an obligation to account for elements which are absolutely beyond the agent's control such as, by their very definition, chance, happenstance, luck. And yet, chance *always* plays a certain role in the success or failure of any *skopos*,

as the Stoics well knew. Thus, this being the case, what conclusion must now be drawn for the status of knowledge?

Despite the limits of Thébert's article, we can nevertheless retain the idea that knowledge should not be seen as some final product of the intellectual virtues, as well as the idea that the notion of knowledge is inseparable from the external influences of chance which affect it. There will always be some part of JTBS that is beyond the control of the agent and, contrary to Greco's definition above ("...it is no accident that *S* achieves those results"), also some part of the agent executing her capacities which depends on accidental luck. Indeed, it was the clear presence of this idea in Thébert's text which inspired me to draw the comparison to the Stoic *skopos*.

But this leads us to a final, considerable problem with Thébert's arguments: due to prioritising, above everything else, the criteria for attributing the status of 'knowledge', she seems to have read the various virtue epistemologies through a lens which blurred as many things as it brought into focus. Virtue epistemology dedicates more of its argumentation to the study of intellectual virtues than it does to what they produce, that is, to beliefs (JTBS, knowledge, etc.). However, the intention of virtue epistemology is not to play the spoilsport to justified and true beliefs, forever looking for opportunities to strip fraudulent knowledge of its honours. Rather, it looks at Gettier cases and situations as limiting cases of our cognitive capacities, the pondering of which enables further improvement of our intellectual virtues, a project which, as we have already seen, could never allow itself to bet on chance. If we take them at their word, then when virtue epistemologists refuse 'gettierized' JTBS as instances of knowledge, they don't do it in order to split hairs over pedantic details, they do it from a desire to sharpen the human capacities of knowing and believing. In short, Thébert has taken a project whose motivations tend towards the *telos* and has studied it in the light of her own project which worries only about the *skopos*. Nevertheless, even though the distinction may seem clear, *telos* and *skopos* must meet somewhere, and it seems as though the border between them must be quite thin, if not even perforated. If we wish to make the application of this Stoic distinction to the Gettier problem a rich and fruitful affair, and not just a convenient one, then we will need a conception of the capacity-to-know and of knowledge which accounts for both the division *and* the relation between its terms. This will be the aim of the next section.

1.3 Knowledge as dynamic manifestation of intellectual virtues

In her reading of virtue epistemology, Thébert sees only a production line with just one function: to create beliefs. Her intention, on the contrary, is to bring out the tighter relation that exists between knowledge and capacity. Like the figure skater's successful quadruple toe loop, knowledge would be essentially and constitutively written into the execution of the capacity and, basing herself on this logic, the author advocates a more supple vision of knowledge, since, in this vision, knowledge *just is* the successful execution of the capacity (*skopos*) and, in that regard, a bit of luck can do no harm... especially to skaters who do not wind up flat out on the ice. But, unfortunately, the comparison that Thébert employs is not entirely fit for the task. Even if, ultimately, she was right to support a conception of knowledge as capacity rather than product, the method she uses in getting there and, as a result, the implications she infers from this, are all inaccurate. Gettier cases are not analogous to the skater who executes her manoeuvre thanks to a spark of luck. Rather, the difference between genuine knowledge and 'gettierized' knowledge is just the same, according to John Turri (2012), as the intuitively understandable difference between the following two situations:

BOIL: You place a cup of water in the microwave and press start. The magnetron generates microwaves that travel into the central compartment, penetrate the water and excite its molecules. Soon the water *boils*.

FIRE: You place a cup of water in the microwave and press start. The magnetron generates microwaves that cause an insufficiently insulated wire in the control circuit to catch fire, which fire deactivates the magnetron and spreads to the central compartment. Soon the water *boils*.²³

This analogy serves to illustrate the fundamental difference between the *manifestation* of a capacity and a consequence for which that capacity is *only* the cause. The first case symbolises the correct workings of intellectual virtues, the second a Gettier case (specifically, we can easily recognise the resemblance it has with **CLOCK**). Let us now examine the subtleties of this analogy.

23. Turri (2012).

The intention behind starting the appliance is identical in both cases, the ultimate consequence for the water also. However, only in the first case can we say that the appliance's proper capacity *accompanies* the realisation of that willed consequence by playing an even more *constitutive* role than was seen in Thébert's examples. In this light, her somewhat belittling attack on the "merely causal" role virtue epistemology attributes to capacity now comes undone. But the analogy does far more than just debunk this critical aspect, as we shall now see as we re-examine Turri's examples using the Stoic distinction from above.

Allowing ourselves an outrageous anthropomorphism, if for the good archer it is a question of *aiming well*, what can we say is the *telos* of the microwave oven? This would seem to be the generation of microwaves capable of rapidly exciting water molecules introduced into its central compartment. Now, as soon as we shift to considering the *skopos*, we must thereby also contend with external circumstances: the *skopos* takes form *around* a *telos* which is unfolding in a particular situation. Therefore, the oven's *skopos* in the analogy is to boil the water, but it could just as easily have been to heat some chicken or to dry some socks. Indeed, whether or not the *skopos* has the actual potential to be realised by means of the *telos* is almost an irrelevant consideration. Ask an archer to employ his skills to topple a castle and he will never succeed. But neither will the endeavour in any way negatively impact upon his capacity to *aim well*. Instances of *skopoi* present themselves to us, we use our respective *teloi* to accomplish them as best we can; sometimes successfully, sometimes less so, sometimes not at all. But, fundamentally, failure is not necessarily due to a mediocre *telos*. Indeed, it is also crucial that the good archer and the good thinker be capable of recognising when the failure of some *skopos* results from external factors or, in fact, from a genuine failure of their *telos*, otherwise how could they ever *improve*? In FIRE, it is an *internal* system problem which causes the failure – a badly insulated wire. In a similar optic, consider the following situation:

BAD ARCHER: A clumsy, war-bound archer grabs his bow and arrows without checking their condition. The wood in the bow has become warped because the archer had stored it in a damp corner. Facing the enemy, he begins to draw his bow, not knowing that his arrow will fly off course due to the warped wood. But as he releases his arrow, he loses his balance and slips in just such a way that this involuntary movement perfectly compensates for the

negative trajectorial effects of the deformed wood. His arrow strikes its unfortunate mark.

It just seems obvious that no army general, having witnessed such a spectacle, would congratulate the archer for this 'successful' act of war. This archer must be reprimanded lest he end up thinking he must find a way of integrating timely loss of balance into his warring 'strategy'. On the other hand, if he has the will and desire to become a good archer, then he will pay little or no heed to this successful *skopos* because he will recognise that something very unusual *and* out of his control was behind it. Should he decide to study the situation, he will certainly discover the deformation in the wood and, in doing so, he will have taken a first step along the path of the good archer. Conversely, had he been congratulated for his successful *skopos*, this praise would have rather had the effect of *distancing* him from this good and perfecting path. Virtue epistemology seems to win out over Thébert's vision of things, practically speaking, for the very reason that it takes this responsibility seriously: setting epistemic agents on the path of the *good thinker* is a *virtuous* objective whose value is obvious.

Concomitant to these somewhat moral ratiocinations, we can also see that the *skopos* resulting from the manifestation of a capacity, like in **BOIL**, is something which *supervenes* on the combination of the *telos* and external factors. These latter can never be fully accounted for by an individual, their variations being potentially infinite. Turri's contention is that we should only speak of knowledge when a belief has resulted from intellectual virtues in a similar fashion; by the manifestation of their genuine disposition. But, as we have already seen in **BOIL**, and contrary to what Thébert read in virtue epistemology, this disposition (*telos*) accompanies the manifestation (*skopos*) *all along its realisation*: it plays a *constitutive* role. If the water is to boil, then the generation of microwaves must continue. Of course the connection is causal, but causality is not reducible to schemas of type "accomplishing *A* caused *B* to commence." To interpret causality in this way is quite simply false. The process of boiling begins with the first microwaves. The process of creating a JTB begins with the first cogitations. In both cases, the processes must continue and accompany the *skopos* so that it can succeed. This continuation is *part of* the respective *teloi* and so, in this way, the manifestation is *dynamic*. If the archer is aiming at a moving target, then he must continuously update his aim. The targets of our intellectual

virtues are likewise often ‘in motion’ or ‘in flux’ and, consequently, we too must constantly update our aim. Precisely, the highest aim of our intellectual virtues must be to constantly adapt to a constantly changing world. In *CLOCK*, had I looked at the same clock again some minutes later and remarked the anomaly of the hands not having moved in the meantime, as a good thinker, I would have, of course, rejected my previous conclusion. I would say to myself that, in fact, as it turns out, *I didn’t know what time it was at all!* I would certainly not tell myself once again that it is eight o’clock! It is in this remark that we can understand the true nature of the *telos*: it is not at all something that can be judged statically, it cannot be reduced to any sum of successful *skopoi*. And it is for this precise reason that this wisdom, borrowed from the ancient Stoics, should be taken on by virtue epistemologists. Virtue epistemology should respond to Gettier cases with a simple, “These are problems with *skopoi* and therefore nothing to lose so much energy over, it happens in all fields. We are not overly concerned with *skopoi*, it is the *telos* of proper thinking that really interests us, since its improvement is something we *can* control.”

Let us now cast our minds back to part of the definition of capacity as provided by Greco: “a stable disposition to achieve certain results under certain conditions.” Having studied Thébert’s criticisms, we have seen that these “certain conditions” are not reliable: unpredictable, changeable and strange situations occur on an everyday basis. Virtue epistemology should therefore take up the Stoic example and accept the omnipresent aspect of chance there is in any manifestation of a disposition, even when the latter is ‘stable’. For the agent, the part that must be identified if she wishes to progress in her thinking, if she wishes to improve her capacities, is just that part which depends on her – her *telos*. Whether her *skopoi* succeed or fail to succeed, she must know, through self-analysis, how to distance the *skopoi* so that she can properly judge her own performance and identify the aspects she *can* and *should* still work on. In the sphere of beliefs, if we say that knowledge is a successful *skopos*, this does not change the fact that any ‘value’ this knowledge itself may have must be disregarded if the most important, most *valuable* thing is seen to be the sharpening of the capacity *to know*. In this Stoic light, the interpretation of the original Gettier case becomes this: Did Smith’s *skopos* successfully obtain? Yes... but this doesn’t matter so much, because his *telos* was an absolute mess. Smith *is* the bad archer, and if he carries on reasoning in such an inadequate way, the boss of the

company will not wait too long before firing him! If virtue epistemology wishes to fully integrate its normative nature, it should adopt something like this Stoic distinction. Yes the world is unusual, yes Gettier cases do occur – as Thébert advances – far more often than we think, but what is the point of fighting over the right to speak of ‘knowledge’ as often as possible when it is in fact the *capacity to know* which is most important to living? In our interactions with the empirical world, we do not have the capacity to attain something more *qualitatively* certain than beliefs since, whether we think of Hume’s merely *regular* vision of the world, of stochastic processes in physics, or of the still only probabilist proofs for the direction in which the arrow of time flies, our best scientific theories reveal to us with ever greater insistence that the targets at which we point our *knowing* are constantly moving. The ‘Stoic’ good thinker must never rest upon his laurels, even in those cases where he has never yet seen his target move or change in appearance, even a target he has successfully struck a thousand times over. He must always remain alert, aware that anything might change at any moment: he must never stop sharpening his *telos*, and this precisely *because* the world is dynamic, unpredictable, and strange.

Yet, on this very point, we must come back to the text of Gettier’s article and to his seemingly anodyne reference to Plato in a footnote on the first page, wherein he asserts that the “justified true belief” definition of knowledge can also be found in the dialogues of this greatest of philosophers: “Plato seems to be considering some such [JTB] definition at *Theaetetus* 201, and perhaps *accepting* one at *Meno* 98” (emphasis mine). Does Plato *accept* the JTB definition of knowledge? Such a conclusion would be difficult to defend! And yet this aspect, this explicit claim within Gettier’s article, has received almost no attention in the vast literature on the subject. So, if I can be permitted to rectify the balance a little on this score, let us now examine the claim.

The Platonic theory of knowledge, clearly expressed in both the *Theaetetus* and the *Meno*, leans first and foremost upon the theory of *anamnesis*, of ‘recollection’, in which, precisely, the soul *recollects* knowledge it had beforehand when it dwelt in *another world* – the world of Ideal Forms. True opinions, Plato states through Socrates, “are not worth much until one ties them down by [giving] an account of the reason why. And that, Meno, my friend, is recollection, as we previously

agreed.”²⁴ For Plato, knowledge is *not of this world*. Any number of references from his corpus – from the allegory of the cavern in the *Republic* to the *Phaedo* dialogue which deals with this question almost exclusively – can be invoked to show that, in Platonic philosophy, true ‘knowledge’ is attained only when one’s life ends in a philosopher’s death, the soul already entirely separated from the body at the latter’s final earthly moment, allowing it to finally leave the sensible world and return to the intelligible one, the world of Ideal Forms. In this sense, Plato agrees with the thesis I have developed here: the accomplishment of a *skopos* always involves an element of chance, for the sensible world (in the Platonic sense, that is, the world in which we dwell, the ‘here below’) is not pure, rather it is a place of constant fluctuation. *Pure knowledge*, on the other hand, can only belong to a world of pure and immutable ideas: Plato’s intelligible world, for instance, or the world of propositional logic exploited by Gettier in the creation of his article’s two situations (an aspect already evoked in footnote #11 above). Unless they be tied down, true opinions “escape,” Plato tells us. Why? Because the reference points that make them *true* are susceptible to change. Certainly, true opinions, once they have been tied down by some account of their causality (an account of their “reason why”), become science,²⁵ but two things must still be made very clear on this point. Firstly, the anchor which transforms true opinion into science is properly situated in a possible, virtual world ‘separate’ to this one. For Plato, no knowledge is possible if it not be anchored in that other world. And, secondly, Plato wastes absolutely no time in stressing that in terms of “directing actions,” that is to say *interactions* with some element of the sensible world, true opinion and science are equivalent.²⁶ Why so? Because of the intervention of this real world’s essence, its mode of being: change, chance, variability. Truth be told, Plato, far from *accepting* the JTB definition of knowledge, was already warning us of the ‘fact’ that any pure knowledge transported to the sensible world, by any means whatsoever, would be incapable of holding on to its epistemic primacy *because the very essence of the sensible world is flux*. Pure knowledge loses its purity *by arriving into the sensible world*. But what do we have here if not a perfect image of what happens in Gettier cases?

Indeed, how could it ever have been thought that Platonic philosophy might “accept” a JTB definition of knowledge when one of Socrates

24. *Meno*, 98a.

25. *Ibid.*

26. *Ibid.*, 98c.

most famous dicta is, "it is likely that neither of us knows anything worthwhile [...] when I do not know, neither do I think I know"?²⁷ Nor should this be understood as some confession from humility: it is the identification of an ontological barrier obstructing the pure existence of any knowledge in the sensible world. Consequently, since Socrates' soul does not yet dwell in that *other* world whose very substance is constituted of pure knowledge (Ideas), he speaks most sincerely in affirming that he *knows* nothing.

But let us leave this Platonic aside for now, in order to come back to it after some further considerations. Keep this in mind though: to affirm that "I *know* something" is to lean on another possible world where it is presumed that no incertitude exists; the intelligible world, the world of pure mathematics, the purely logical world, or else, of course, some world of one's own creation, where all of one's personal beliefs *are* certain (since the individual in question created it *for* this to be the case) and where these beliefs are, therefore, for the individual inhabiting that world, instances of certain *knowledge*. This final reflection of this first part will have a fundamental importance for the analysis which is to follow in the second part.

All of the preceding argumentation represents only half of the work to be done here. It is simply not satisfactory to content oneself with a purely theoretical approach when tackling a subject as important and concrete as the ideal workings of the human mind. If the real world really is such a labyrinthine epistemic flux, then we must go beyond the sphere of imaginary examples and confront this real-life world head on. Consequently, the second section of our analysis will be given over to examples taken directly from our strange everyday world and to mankind's observed cognitive behaviour when he finds himself in certain common situations: behaviour, nevertheless, which flies in the face of our natural intuitions about the workings of a 'normal' mind.

27. *Apology*, 21d. For consistency, I have taken all quotations from the same Cooper & Hutchinson edition of Plato's Complete Works. Other translations, however, make Socrates' dismissal of the possibility of certain knowledge far more explicit, for example; "I, as I do not know anything, so I do not fancy I do."

Resolution of Cognitive Dissonance

The surprising priorities and motivations of the cognitive agent

MY GOAL NOW is to observe and examine beliefs in their natural environment: the mind of the cognitive agent. Some interesting discoveries from psychology have revealed new insights into the importance an agent accords to her beliefs when it comes to their reviewal and revision.²⁸ The implications of these discoveries, already widely discussed in philosophy (*e.g.* Harman 1986, Lehrer 1990), deserve to be employed in the study of Gettier cases. Significantly, if the cognitive agent, during certain psychological experiences, does seem to arbitrarily upgrade or downgrade both truth value and justification when it comes to choosing between conflicting beliefs she may hold, then we should ask ourselves what reasons could explain this apparent epistemic anomaly and what could explain the fact that this same anomaly seems to occur on a recurrent, spontaneous, and practically universal basis. Hence, the first task in this second section, once the theory of cognitive dissonance has been outlined, will be to study Leon Festinger's classic experiment, conceived to demonstrate the pertinence, naturalness, and spontaneity of cognitive dissonance resolution through revision of beliefs. Revisiting the Platonic notion of a world of Ideal Forms, I will show the family resemblance between the general resolution mechanisms set in motion by cognitive dissonance and the specific ones employed in the creation of Gettier cases. On the basis of this comparison, coherentism will emerge as the philosophy which best corresponds to the cognitive agent's actual workings. It is Gilbert Harman's (1986) take on coherentism that I will put to the test, for the reason that it makes use of very similar situations to those conceived of by Festinger in his experiments. This testing of coherentism will rely in part on the philosophy of Jean-Paul Sartre (1943, 1945, 1946) and will bring us back to the 'Stoic' virtue epistemology sketched out in the first section. Finally, having held our analyses of Gettier's hypothetical cases up against our analyses of Festinger's real life experiments, we will see that, while the proposed new virtue epistemology did constitute a step in the right direction, it didn't go quite far enough in its prescriptions:

28. There are dozens of such sources. I will mainly be concentrating on Festinger (1959) and on Anderson, Lepper & Ross (1980).

the *skopos/telos* distinction will have to be shored up by an additional demand of ‘authenticity’ which, alone, can open our eyes to the unconscious refuge we take in our virtual worlds and thus enable us to put this same capacity to good use.

2.1 Cognitive Dissonance – the spontaneous relegation of JTBs

In his renowned book *A Theory of Cognitive Dissonance* (1957) Leon Festinger, searching for experimental proof to support the new ‘cognitive dissonance’ hypothesis he had conceived, studies the results of several dozen psychological and sociological studies related to the cognitive mechanisms and influences that condition belief revision. A colleague of Festinger’s provides us with a succinct definition of the concept: “Dissonance is a negative drive state which occurs whenever an individual simultaneously holds two cognitions (ideas, beliefs, opinions) which are psychologically inconsistent. Stated differently, two cognitions are dissonant if, considering these two cognitions alone, the opposite one follows from the other.”²⁹ The classic illustration of this phenomenon is that of the smoker who wholly believes that smoking causes fatal cancers but who, nonetheless, continues to smoke. Thus, taking up the preceding definition, and presuming that the individual does not want to die of cancer, the action consisting of *stopping* smoking should follow from the belief that smoking kills. But, does this imply that the smoker himself considers that he is nurturing psychological incompatibilities if he not stop smoking? No. Or at least, not necessarily. The normal psychological response to this kind of situation consists of seeking out rationalisations for the dissonant activity. Our smoker has two incompatible mental contents and, therefore, by assumption, he will experience cognitive dissonance. In order to resolve this dissonance, the smoker has two options: either stop smoking or else seek out mental contents which are *consonant* with both the fact that he smokes despite knowing it is fatal and the fact that he doesn’t wish to accelerate his own death. The principal difficulty with this choice is that the first option is subject to a considerable obstacle: it is not easy to stop smoking. With respect to the second option, however, it is very easy to find a whole host of strategies and arguments which, in a far less burdensome way than the first option, will render the dual act of smoking-despite-being-aware-of-its-dangers *consonant*. Festinger gives several examples.

29. Aronson (1969), p. 2.

The smoker can tell himself that, "(a) he enjoys smoking so much it is worth it; (b) the chances of his health suffering are not as serious as some would make out; (c) he can't always avoid every possible dangerous contingency and still live; and (d) perhaps even if he stopped smoking he would put on weight which is equally bad for his health. So, continuing to smoke is, after all, consistent with his ideas about smoking."³⁰ There is nothing new about this observed phenomenon of seeking out justifications for one's deeply rooted beliefs and avoiding or discrediting information that challenges them. The novelty Festinger brought to the table was to explain this by means of an unpleasant psychological tension, a pressure, a negative drive to *retrieve* a state of consonance, in place of what had been the standard view in psychology up until then, *viz.* a positive process of egocentric reinforcement.³¹ In a way, Festinger defused this rather psychoanalytic conception in order to replace it with the uncontroversial explanation of a universal individual drive towards cognitive consonance; the creation of a coherent basic mental state, a default position free of intrinsic conflict, necessitating minimal psychic maintenance, subsisting almost autonomously, a kind of psychological haven through the windows of which the agent can gauge all novel information about the world she experiences. This new information can then be accepted or rejected in its current state, or, as in the case of the smoker, it can be painted over, so to speak, with other beliefs that render the new information anodyne: "Yes, I know that smoking kills, but I just enjoy it so much! And also, the way I see it, we're all going to die one way or another, why not spoil ourselves a little while we wait!" Indeed, the example shows just how easy and banal it is to defuse a conflict and create a coherency out of two opposing beliefs, specifically, "I don't want to die" and "This activity is fatal to me." One simply paints over the second belief and, as if by magic, not only does the smoker still not want to die but, thanks to this newly dressed up belief, it is in fact *by smoking* that he will show just how much he wants to live to the fullest! Of course, the example is mere anecdote, but still it illustrates perfectly this capacity for creative and subjective re-interpretation that each of us surely recognises, be that in ourselves (if we are honest) or only in others (if we are not). Is it not, in fact, this self same capacity we refer to in the expression 'sour grapes'? An expression inherited from this famous fable...

30. Festinger (1957), p. 2.

31. In psychology, this type of process is called *self-serving motivation*, *egocentric bias*, or *egocentric attribution*. Cf. Heider (1958), Ross & Anderson (1982).

*A fox from Gascoigne, or perhaps Normandy, did traipse,
 And, nigh-on starving, did spot high on a vine
 What looked to be some ripened grapes,
 Skin already the colour of wine.
 The lad would have made o' them a tidy dish;
 Yet, alas! his reach was a little too short:
 "They're too sour!" he decided "Fit only for fish."
 Well, better than complaining, to react of the sort!³²*

When we recall that La Fontaine borrowed this fable from an original attributed to Aesop (620-564 BCE), the least we can say is that the strategies for *resolving* cognitive dissonance have been observed by humankind for a very long time, albeit without receiving an official label.

Festinger's 1957 book is a case by case identification and interpretation of instances of cognitive dissonance in psychological studies carried out by other psychologists. Subsequently, having managed to successfully apply the hypothesis to explain a wide variety of diverse behaviours, Festinger and his colleague James Carlsmith decided, in 1959, to conduct their own experiment, the objective of which would be to clearly demonstrate the workings of this extremely widespread (if not universal) cognitive mechanism.³³ I will now describe the essential elements of this experiment before we then examine it through the lens of the conclusions from the first section.

The core of the experiment consists of intentionally creating an incoherence between a genuine private opinion and a faked public opinion. Three subject groups are formed, made up, in the original experiment, of psychology students. All subjects are made to undertake an hour long task. The task in question is designed to be easy, very repetitive, and very boring, the idea being to have each subject form more or less the same private opinion of it: the task is not fun. Having completed the task, the first two groups are then led to believe that the experiment's actual objective is to measure comparative performances between subjects who, like them, have no prior opinion of the task to be undertaken and other subjects who have been given a positive prior opinion of it before commencing (the third group being the control). Once the two groups have been let in on the 'secret', they are then individually asked to communicate this positive opinion ("the task is

32. Jean de La Fontaine, 'Le Renard et les Raisins'. Translation mine.

33. Festinger & Carlsmith (1959).

really fun, interesting," etc.) to the next subject because, they are told, the laboratory assistant who would normally do this has taken ill. The experimenter apologises for the unusual nature of the request and offers \$1 to each subject in the first group and \$20³⁴ to each subject in the second group as compensation for, essentially, peddling this little white lie that the task is fun (though, of course, it is not presented in these terms). The subjects thus prepared, they each spend 5 minutes talking to another subject (who, in this instance, really is an assistant to the experimenters) about the experiment, expressing only positive things. Following this, they are interviewed about their own opinion of the experiment in a separate room and by another assistant.³⁵

Before analysing the results produced, let us look at Festinger's own predictions by quoting from the abstract of the experimental report: "1. If a person is induced to do or say something which is contrary to his private opinion, there will be a tendency for him to change his opinion so as to bring it into correspondence with what he has done or said. 2. The larger the pressure used to elicit the overt behavior (beyond the minimum needed to elicit it) the weaker will be the above-mentioned tendency."³⁶ In other words, the less the subject has to be *pushed* to act contrary to her private opinion, the more she will end up identifying with or internalising the implicit opinion attached to this 'dissonant' action. Thus, \$20 being a considerably greater recompense than \$1, the expectation was that the subjects in the \$1 group would have a higher tendency than those in the \$20 group to internalise the false opinions they had just verbally communicated. And, indeed, this is precisely what was seen to happen. Following the session of persuasion, the \$1 group answered much more positively when asked for their own opinion of the initial task when compared to both the \$20 group *and* to the control group who had done no persuasion. The subjects from the \$1 group gave

34. The equivalent today to more or less \$10 and \$200.

35. For reasons of space, I must leave out many of the detailed precautions Festinger and Carlsmith put in place in order to avoid any suspicion on the part of the subjects that they may in fact have been participating in an experiment with a totally different objective, as was indeed the case.

36. Festinger & Carlsmith (1959). The same predictions are already to be found in Festinger (1957), pp. 84-97.

the task much better marks and, furthermore, were more inclined to accept to do the experiment again should the need arise.³⁷

The idea behind this experiment, and others like it, is to create a situation which will tend to produce an *objectively* or *universally* held opinion, for instance, in this example, that the task is boring and monotonous. In this way, any private opinion pointing in that direction can be said to be more or less objectively 'true'. The fact of having actually lived through the experiment then 'justifies' the opinion, such that we can call this private opinion – that the task is boring – a (non-gettierized) justified and true belief; an instance of 'knowledge'. But what about the opinion that is held *after* the five minute persuasion session? What is *its* epistemic status? Firstly, the subjects have been lying, so it cannot be true. Secondly, it involves replacing an empirical justification (first-hand experience) for an *ad hoc* justification which imposes itself after the fact (the somewhat complicated nature of which we will discuss presently). Finally, it is strictly speaking difficult to even call this second opinion a belief at all, no more than an actor playing Hamlet can be said to believe that he himself has actually killed Polonius. And yet, the results of the experiment show that in these quite normal experimental circumstances (normal for psychology students in any case), the private opinion (a JTB) is abandoned in favour of some sort of chimerical mental content which we hesitate to even call a belief. How... *why* would certain cognitive agents internalise a belief that is both built out of a lie and in contradiction with their own direct experiences? A clue is to be found in the explanation as to why the phenomenon has a higher occurrence rate in the \$1 group than in the \$20 group: the less the subject feels *pushed* to act contrary to her private opinion, the more she will end up identifying with or internalising this action. The action in question here is a speech act and, by internalising it, the subjects simultaneously internalise its propositional content. With the \$20 group, the subjects have sufficient *external* cause to pin an explanation for their action on, an action which

37. It should be mentioned here that Festinger recorded these discussions in order to have them studied by an independent panel. This was to verify that the \$1 group hadn't simply been more convincing than the \$20 group, convincing themselves by the strength of their own ingenious arguments and thus independently of the internalising effects of cognitive dissonance resolution. The panel found no significant difference between the two groups' arguments. This 'self-persuasion' argument has been dismissed several times (*cf.* Festinger, 1957). It should also be noted that this experiment is only the first in what was to be a multitude of similar experiments – the same predictions have been supported numerous times through various other experimental methods.

consequently does not modify the rest of their general beliefs about the world and their place in it. These subjects leave the persuasion session without experiencing dissonance, something which is not the case for the \$1 group. For the latter, the fact of having expressed an opinion contrary to their own in exchange for very small recompense is dissonant with their global view of themselves and their general beliefs about the world. In carrying out this action, a dissonance has arisen and created a tension that the mind must now try to resolve by reviewing its vision of the world. But what is the underpinning nature of this vision?

The easiest way to think of it is as a network, a set of interconnected beliefs. Our beliefs are not independent, they act upon each other, they push and pull each other, support or contradict each other. Recall the situation **ROBOT** from the first section: Mary believes that the robot identical to her husband *is* her husband. But if one fine day Mary is eating dinner with her robo-husband and, lo and behold, her real husband comes knocking at the door, Mary will have to immediately review her vision of the world. She cannot simply say, "Oh look, I actually have *two* husbands!" Moreover, she would find herself in a situation of acute cognitive dissonance. Precisely what Festinger invites us to conceive of is that, already *during* the persuasion session, the conflict of beliefs consisting in *believing* that "the task was boring" and *saying* that "the task is really interesting" gives rise to an explanation in the \$1 group subject's mind that is as worrisome to her as it is inevitable, and which takes form even as she is in the act of articulating the false opinion. That emergent explanation, simply, is this: "I compromise on my true beliefs for very little!"

In order to escape this situation as efficiently as possible, let us imagine that the mind operates according to certain economical laws, laws of least effort. If this is the case, then the mind's first recourse will be an attempt to modify those elements which offer the least resistance to change: "To the extent that [an] element is consonant with a large number of other elements and to the extent that changing it would replace these consonances by dissonances, the element will be resistant to change."³⁸ Therefore, among the beliefs implicated by the dissonant situation, it would be a case of choosing the least 'connected' belief in a cognitive network which is otherwise globally coherent, modifying *only* that belief and then surreptitiously sliding it, in its newly altered form, back into the network – defused, anodyne, consonant.

38. Festinger (1957), p. 27.

This image of beliefs as a network should remind us of something. After all, what does a global network of beliefs constitute if not a certain *possible* world of ideas? Except that, inside the mind of a given cognitive agent, this world is not subjectively experienced as a construction built of *beliefs* but of *knowledge*. On this, and speaking more as a philosopher than a psychologist, Festinger perspicaciously states that, “An individual does not hold a belief unless he considers it to be correct, and thus, psychologically, it is no different from ‘knowledge’.”³⁹ But let’s not stop there; if, from the agent’s subjective perspective, these beliefs do indeed take on the solidity of knowledge, then this is precisely because they *do* possess real structural force and effective function – these beliefs hold up an entire world! Hence ‘knowledge’, *psychologically speaking*, is just a status the cognitive agent attributes to her own beliefs with respect to the coherency and structural stability she (correctly) observes in them.

Basing myself on this conception, I would like to propose an analogy that illustrates the prolonged effort needed to create such a world, the reticence liable to arise should it need modifying, and the choices one makes regarding necessary modifications which may nevertheless impose themselves.

THE PAINTER AND THE WHITE LIE: A painter toils for years to create an enormous painting of an imagined landscape, a vast stretch of bucolic countryside, mountains and forests in the background, nothing but unadulterated nature all the way to the horizon. So far as possible, the artist’s vision is for this painting to faithfully reflect his own spirit as he contemplates nature in all its glory. But one day, painting in his meditative, automatic way, his thoughts stray to the unpleasant weekend he has just spent with his nephews in the city. Ushering these memories away, he casts his attention to the canvas and notices that, in his troubled reverie, he has begun to paint a skyscraper where his fields meet the mountains. Its presence is completely incoherent with the rest of the landscape! And yet, it has arisen as a spontaneous reflection of his spirit in that moment! What should he do? If he leaves the concrete eyesore, he will remain true to his artistic vision of faithfully representing his own

39. Festinger (1957), p. 10.

spirit, but he will also forever have to explain to one and all why there is an unfinished, ugly building right in the middle of his beautiful painting. Two options for making the painting coherent again present themselves: either he can transform the entire canvas into a representation of the city, from now on saying that this is what his spirit was meditating on. Or else he can paint a couple of trees over the building and never mention his little slip of the mind and brush to anybody.

Intuitively, one would say that the painter really only has one *sensible* choice. The other choice implies an enormous expenditure of energy, and for what result? Just to account for a tiny stray thought which, through temporary carelessness, ended up on the canvas. It would be absurd to follow this option and, indeed, this imagined quandary seems silly in the extreme. Yet, in reality, the \$1 group are faced with exactly the same kind of massively unbalanced choice: the boring accomplishment of the task during the experiment is nothing but a stray moment when seen through the global lens of all the episodes that make up an individual's life. Internalising the white lie "I liked the task" costs virtually nothing in psychic energy. The subject can easily revise her original belief in this way without impinging on other more important beliefs. In short, in order for the agent's world of ideas (her network of beliefs regarding the world and her place in it) to retain its structural integrity, a small true belief must be replaced by a small *fictional* belief. If we imagine this world of ideas like a house, then there is a veritable hierarchy of importance at stake: there are beliefs which constitute the foundations, the beams, the walls – load-bearing, *supporting* elements. But on the other end of the spectrum, there are also beliefs that are more like the bottle of vodka in the cupboard or the bit of string in the drawer – *incidental* elements. These latter beliefs lend themselves more easily, more imperceptibly, and, above all, more *independently* to modification. In this perspective, the subjects from the \$1 group did nothing more threatening to the overall integrity of the house than to drink just a little mouthful from the bottle of vodka, fill it back up to the correct level with a few drops of water, and slip it back into its original place. The only problem for the experimental subjects above is that, unbeknownst to them, their vodka pilfering was being observed!

This process of creating and maintaining a mental world, in philosophical terms, is a creative process of *rationalisation*. Indeed,

Elliott Aronson speaks quite accurately when he says “dissonance theory does not rest upon the assumption that man is a *rational* animal; rather, it suggests that man is a *rationalising* animal – that he attempts to appear rational both to others and to himself.”⁴⁰

What then is our conclusion at the end of this analysis of Festinger’s experiments? The mind of each individual is now seen as upholding a network of interconnected beliefs which forever seeks consonance, coherency. At any given instant, this network can be quite literally considered as a self-contained, coherent world of ideas. Should ever it seem that the empirical world the mind is normally engaged to represent is becoming *less* coherent than an individual’s already mentally constructed world of ideas, then the mind can and may well choose to ignore a given consonance threatening empirical truth and, in acting thus, reveal that, in fact, it prioritises coherency *above* truth. On the basis of the earlier reference to Plato, I suggest that it is *because* the mind is constantly exposed to the flux of the empirical world that it acts in this way when confronted with a conflictual trade-off between ‘subjective coherency’ and ‘objective truth’.

The most important result from this analysis, however, is that the new mental content “the task was very interesting” now appears to be as much an instance of ‘knowledge’, within the strict world-of-ideas limits of a subject from the S1 group, as did the proposition “the man who will get the job has ten coins in his pocket” within the limits of the manoeuvres allowed in the world of propositional logic, as exploited by Gettier. Our intuitive suspicion of something going awry when these two propositions are labelled as ‘knowledge’ stems from the fact that what we are dealing with, in both cases equally, is the execution of certain operations within some *other* world followed by an attempt to bring the epistemic identities created according to *that* world’s rules into line with the normal, empirical rules of *this* world. Simply put, the laws of propositional logic, the laws of psychological consonance, and the laws of nature are not the same laws. It would be as if our painter, having succeeded in coherently uniting upon his canvas a horn and a horse in the form of a unicorn were then to be surprised to discover the inability of a veterinary surgeon to reproduce the same success in flesh and bone. The painter passes by the world of his imagination to succeed in his depiction, Gettier passed through the world of propositional logic

40. Aronson (1969), p. 3.

for the same reason. As for the cognitive agent, she takes refuge in the consonance of her own virtual world rather than confront the complex empirical truth of certain situations she deems to be, literally, not worth the effort.

Festinger's experiments seem to show that the empirical *truth* of our beliefs is a less important notion in their constitution than is generally admitted and that, in reality, rather than aiming to ground our beliefs first and foremost in empirical experience, we seem needful of casting them in such a way as they will harmonise with the network of all our other beliefs, this network being none other than the world of *our* ideas, each of us, individually. And does this thesis not provide a rather interesting explanation as to why, in empirical situations which prove to be too traumatic, too stressful, too absurd, simply too psychologically difficult to assimilate, certain individuals retreat into what we call a fantasy world? Fantasy, yes, but, most importantly, internally self-coherent.⁴¹ Most often, even the individual deemed to be mentally insane still remains a *rationalising* animal who bends empirical reality to coherently fit into whatever mental, virtual worlds they happen to be occupying psychologically. The problem for the 'insane' mind, what makes it *remarkable* among all other socially conforming minds, is not the fact that it acts independently of *the empirical world* (because, precisely as cognitive dissonance resolution shows, we *all* do this), but that it acts independently of *the socially established norm*.

Consequently, the strategies underpinning the resolution of cognitive dissonance seem to be purely coherentist. However, must this necessarily lead to the inference that a coherentist approach to epistemology is the best approach to adopt? In order to answer this question, let us now turn our attention to another strong version of coherentism within epistemology and see just how far this philosophy can bring us in providing philosophical answers to the experimental data on cognitive dissonance and the strategies employed by the cognitive agent in avoiding or resolving it.

41. On this subject, see also the reverse reaction, for instance with Paris Syndrome where certain Japanese tourists, having an almost fairytale conception of Paris before visiting, undergo genuine psychological crises when they are confronted with the somewhat less idyllic reality of the living, breathing, modern metropolis that it is *in fact* (Viala *et al.*, 2004).

2.2 Coherentism or “bad faith”?

In the introduction to this second section, I said that Gilbert Harman – in the chapter ‘Positive versus negative undermining’ of his book *Change In View* (1986) – tackles very similar situations to those analysed by Festinger. It is, therefore, interesting as a starting point to note that Harman makes no mention of the famous psychologist or of his even more famous theory. Ross & Anderson (1982), whose article Harman draws from quite extensively, do not, however, commit the same fault of omission, disqualifying any potential claim of ignorance on Harman’s part. One reason which could, however, explain this omission is that Harman’s own recurrent thesis, the principle of “clutter avoidance,” fits quite poorly with cognitive dissonance theory. Indeed, to give a global idea of the difference between the two approaches, where Festinger explains a set of psychological phenomena on the basis of a specific cognitive *dissonance* mechanism, Harman postulates around cognitive *limits* in a quite general sense. The clutter avoidance thesis revolves entirely around the mind’s finite capacities for memory and memory retrieval: the agent must avoid collecting gratuitous mental contents and save only what is, in a broad sense, relevant. Let us briefly consider the example Harman uses in *Change In View* and then look to various interpretations of the phenomenon it illustrates, whereby beliefs persevere beyond the point where their foundations have been explicitly discredited.

Harman tells us the story of ‘Karen’, who receives somewhat surprising results from a recently taken aptitude test. According to the results, she has a strong aptitude for science and music but a weak aptitude for history and philosophy. Prior to this, however, she had succeeded rather well not just in science but also in history, whereas she didn’t do quite so well in music. In short, half of her previous aptitudes seem to have been reversed, according to these latest results. In keeping with Aronson’s slogan, Karen proceeds to a process of *rationalising* this new information. Harman tells us that, having rationalised the new results via inferences cast through the lens of certain relevant memories (for example, “In fact, thinking about it now, my old history course really was too easy,” etc.), she assimilates and internalises the new results and, on this basis, modifies her future study projects: she will, she decides, abandon history in order to pursue music.

But we are then further invited to imagine that, a few days later, Karen receives a call telling her that those results were not actually

hers and that, unfortunately, her actual results have gone missing. The question which then presents itself is this: what should Karen now do with the change of project she has given herself in light of the (as she now knows) erroneous results? Harman states that the foundationalist response fits with our natural intuitions by stating that she should abandon the changes she has made because their foundation (the test results) has now been wholly discredited. The coherentist response, on the other hand, takes the overall coherency of her reasoning into account and sees enough justification in it for her to continue pursuing these new decisions. However, says Harman, in contrast to our elementary intuitions, it is in fact the second *dénouement* which occurs in reality: people, he claims, generally do not reject the results of their reasoning *even if* the foundations of that reasoning have been swept away.

Harman then goes on to relate his illustration to analyses conducted by Ross & Anderson (1982) into the perseverance of beliefs formed by subjects of psychological experiments even beyond debriefing sessions during which the fictitious nature of the experiment's entire basis of information and results is revealed. The consequence of such a revelation for the subject *should* be the realisation that their supposed 'good' or 'bad' performance report has, in fact, no value whatsoever. The specific experiments in question involved, in one case, sorting out genuine from fictitious suicide letters and, in another, drawing psychological links between the risk-taking preferences given by a firefighter and his actual performance in work (*i.e.* when faced with a blazing fire). The subsequent debriefing sessions consisted of informing the subjects that all the suicide notes were in fact fake and that the entire collection of firefighter risk-preference versus performance information was pure invention. And yet it was found that, when questioned after this debriefing, the subjects still persisted in their positive or negative estimation of their own performance, an estimation whose sole apparent foundation was results which they must now have recognised as being pure experimental fabrication. So, what interpretations can be drawn as to what may underpin this phenomenon of belief perseverance?

Harman maintains that, in general, the agent does not *keep track of* her justifications and, given that everything seems to be coherent within her *current* mental state, she does not in fact realise that the unique reason the belief in question had for existing no longer has any value whatsoever. For Harman, the agent has simply "forgotten" the primordial role played by the founding justification (the 'fake' results of the tests): "[The subjects] continue to believe things after the evidence

for them has been discredited because they do not realize what they are doing. They do not understand that the discredited evidence was the *sole* reason why they believe as they do.”⁴²

In Festinger, the explanation for the same set of phenomena would consist of showing that, while the specific nature of the tests does not have much importance in the everyday life of the subjects and will thus be incapable of provoking any kind of considerable consonance or dissonance, the underlying faculty of judgement, on the other hand, is something whose perceived proper functioning has fundamental value for any individual. Having mobilised an entire set of memories and ratiocinations in order to arrive at her judgements, if an agent were then to try to mix in the new mental content “my judgements on this subject were worthless,” then it would be *this* self-avowal that would provoke significant dissonance. The agent, therefore, will convince herself that, even though this specific, piddling little experiment was just experimental trickery, she knows that her *reasoning* would have been good (or bad, depending on what the final result of her self-assessment had been) in a genuine situation. Indeed, in their 1982 article, Ross & Anderson refer to a second version of the experiment (Anderson *et al.*, 1980), again using firefighter performance evaluations, but in which the subjects had to *explain* the reasons why they had come to such and such a conclusion regarding each firefighter’s psychological make up. In these cases, the belief perseverance was distinctly stronger following debriefing than it was where no such explanation was demanded. In other words, the more the agent’s global and fundamental faculty of judgement was called into action the more the agent tended to hold on to her beliefs after the epistemic content of the experiment had been fully discredited. Or to state it more plainly still, even in cases where the self-assessment leads to a negative estimation of the agent’s own psychological discernment, the tendency is still to conclude negatively on *that* aspect rather than to admit that her most fundamental mental faculty was flawed: her very capacity for rationalisation itself.

Ross & Anderson themselves evoke a process of recollection by which the agent compares her success or failure in the current experiment (for example, let us imagine the case where the experimenter informs the subject that she has utterly failed to correctly sort out the suicide letters) with other similar successes or failures from the past (for example, “I got a really low result in a psychology test I did in a magazine once”). After

42. Harman (1986), p. 38.

the debriefing, *all* of these similar successes or failures are still present in working memory and therefore strongly influence the agent's self-evaluation following the 'fake' experiment.⁴³

In Harman's interpretation, we see clearly to what extent he insists on the *limits* of subjects' cognitive capacities, to a degree that could almost be seen as insulting. And yet, the wager he makes in betting everything on the supposed fact that the subject simply forgets justifications seems very difficult to defend. His intent is to reject the foundationalist response whatever the cost and, as a result, he doesn't see that the mobilisation of similar, older life experiences (which root the agent's *new* cognitive situation in an older, more global and coherent set) creates a genuine, broad *foundation* for these beliefs whose relevance is much deeper and more significant to the agent than the relatively minuscule foundation upon which relied only the trivial content of the specific psychological experiment at hand. The global foundations of the agent's faculty of judgement and its successes or failures, anchored in place by a whole life's worth of experiences, work to assimilate the novel experience into a coherent network. The agent loses track of nothing, she *replaces* the discredited evidence with older and, therefore, for her, more concrete evidence. Whether the evidence is negative (a set of beliefs of type "I have often misjudged emotional situations") or positive (of type "my friends often come to me to talk when they are distressed"), the agent, in answering the questionnaire following the debriefing, manifests perseverance of these beliefs for the very reason that she answers on behalf of her *general faculty of judgement*, to which has just been added, *after the fact*, the specific details of the psychological experiment which, for *her*, for her life-encompassing *coherent identity*, hold virtually no importance and, quite possibly, have acted only as a trigger for a fresh reviewal of old information. This is quite precisely the same cognitive activity demonstrated by Festinger. Likewise, it should be quite clear that Karen, in Harman's illustration, would not have altered her projects on the basis of the erroneous aptitude test results were she genuinely attached to her identity as "someone who is good at history"; if this had been the case, she would have been more likely to question the test itself rather than her own academic capacities! It is difficult, for example, to imagine an Oxford don suddenly abandoning history because of a failed aptitude test! And this is exactly the direction in which Ross

43. Cf. Ross & Anderson (1982), pp. 150-151.

& Anderson's interpretation goes, which is precisely why it just seems wrong or misdirected to turn to limits in memory in order to posit that the agent just stupidly forgets the justifications for her beliefs and that, *therefore*, foundationalism can not be true.

Festinger's theory shows that the cognitive agent chooses subjective consonance over objective truth in cases where the latter threatens the former's stability. But Harman's coherentism gives us something completely different; *his* agents are obliged to be coherentist because they don't have the mental capacity to keep track of truth. In order to be epistemically coherentist, Harman posits that the empirical world does indeed contain fixed and reliable truths, the problem for him is that man is incapable of successfully tracking them. Ultimately, Harman seems to criticise humankind for not succeeding in its epistemic *skopoi*. Clearly, this vision cannot fit with a more Humean view of a fluctuating empirical world which the agent experiences but decides to retreat from by leaning on her capacity for inference, her capacity to create virtual and coherent worlds which provide her with more solid foundations than those available in the empirical world – more solid because it is the agent herself who *decides* upon their solidity even as she is creating them (whether this be a conscious or unconscious 'decision' – it is, in either case, a decision making action that she can *become* aware of).

The view the cognitive agent forms of the sensible world is not completely fluid. She retreats into fixed beliefs – about herself, about her entourage, about her world. She puts up with *small* changes, but large ones stress her, may even drive her crazy, lead her to reject data coming directly from the sensible world. She would like to believe that the sensible world is fixed and immutable, what in ontology would be called an *in-itself* or *an-sich* world, having an assured and reliable essence just waiting to be discovered. And this is what Harman seems to be suggesting is possible; that if man only had the enormous mental capacity necessary to keep track of all of his justifications then he could, potentially, hope to pin down the fundamental, fixed nature of the world, its essence.

In contrast, Festinger's cognitive agent is clearly an embodiment of the Sartrean ontological paradox of "*not* being what one is" and "*being* what one is not"; the problem in this view arises not from man's finite memory but rather from the fact that he does not take the responsibility in his choices that would befit his *infinite liberty*. A thinker like Sartre would simply say that the subjects from Festinger's \$1 group are acting out of "bad faith," would say that they are simply lying to themselves,

taking themselves, even in the act of their lying, to *be* honest. True honesty, however, as with most other character traits, is no essence one can *have*, it arises in the free *acts* through which individuals manifest themselves and *only* in these acts. Indeed, it is in this very sense that Sartre describes Proust's genius: "it is neither the oeuvre considered in isolation, nor the subjective ability to produce it: it is the oeuvre considered as the whole set of the person's manifestings."⁴⁴

Thus, taking back up John Turri's vocabulary, to always work on this manifesting and avoid resting on the accomplishment of past manifestations (whether this be manifestations of genius, of honesty, of intelligence, etc.) is equivalent to always working on one's *telos* independently of the failure or success of any related *skopoi*. Only a conduct of "bad faith" can explain the behaviour of the subjects in the psychological experiments we have seen, and the same conclusion applies in the tale of Karen and her aptitude tests. If these individuals were *authentic*, in Sartre's sense, then they would not have tried to hide the liberty of their choices (the choice to lie, the choice to adopt a false belief so as not to appear to themselves to be liars, the choice to change career projects, etc.). This is because, to be authentic is to be authentic in relation to a *project*; a permanent *willing* or *telos* which personifies, manifests, and accompanies the individual agent in all of her actions and all of her choices. If Karen was authentically living out such a project, then the results of some test, whether they be erroneous or not, could not have distracted her from this authentic will. Yes indeed, her behaviour – and that of the subjects in the experiments – seems to be completely normal and widespread, but then it is just obvious that any project aiming to *improve* intellectual virtues must begin by studying what is common sense in order to identify what can and 'should' be improved. A normative discipline, like Stoicism or virtue epistemology, must identify and describe normal, common sense functioning *in order to* distinguish itself from it. Tying knowledge and knowing to the common sense *skopos* and Stoic *telos*, respectively, enables us to grasp why it is that the agent's aim should only be to *perfect*, to *always be perfecting* the *telos*. Likewise, tying the resolution of cognitive dissonance to the conduct of "bad faith" enables us to grasp why one should aim for a conduct of authenticity which becomes vulcanised *around* such a *telos*.

Further exploring this Sartrean analysis of belief perseverance and resolution of cognitive dissonance would necessitate a deep investigation

44. Sartre (1943), p. 12. Translation mine.

into the French author's entire moral philosophy. But, within the limited framework of the present study, it can simply be observed that, once again, assiduously tracking the *believing* that is actually employed by the agent, as well as the believing that she *should* be employing, leads to certain considerations about her moral responsibilities and brings us back again, albeit by another route, to the necessity within epistemology of a strong normative discipline such as virtue epistemology. Finding ourselves at this same point again, we can now conclude by sketching out an outline of what this new *authentic* virtue epistemology might consist of.

Conclusion

Towards authenticity in epistemology

MAN IS a *rationalising* animal and the manifestation of this rationalisation is the creation of virtual worlds which are rational and coherent *for* the sensible world, *on its behalf, when it is not*. If man does risk forgetting something, then it is not so much the justification of his beliefs, as Harman maintains, as the fact that it is he himself who creates the conditions for *all* justification, who creates the notion of truth in his world of Ideas before superposing it and imposing it upon the real, sensible world. This is the exact sense in which Sartre asserted that “man is the being by whom truth appeared in the world, his task is to give himself entirely over to turning the natural order of existent entities into an order of truths.”⁴⁵ The same Hume quoted by Angélique Thébert could also remark that man has adopted the *habit* of creating worlds which are only *possible* (through inference) to such an extent that he no longer pays any attention to this act, he often does it without thinking, expecting the sensible world to conform itself to one of his virtual worlds. After all, if that doesn’t work, it’s not the end of *all* worlds; he can always create another one.

The consequences of this absolutely commonplace habit of confusing the content of possible worlds for the content of the empirical one are to be found, in more or less distinct forms, in both Gettier cases and in the psychological strategies employed in resolving cognitive dissonance. This does not at all make a ‘false’ problem of Gettier’s titular question. Rather, it reframes it within the full scope of what it actually implies: the definition of knowledge problem just *is* the problem of defining our human cognitive nature. Conversely, responding to Gettier with nothing more than a new definition for knowledge is failing to understand that what is at stake *is* a genuine problem of metaphysical limitation whose answer must consist in analysing the origin of this limit and suggesting methods for how man should act having now admitted to himself that it actually exists. Sartre gives an example of such an answer in the man who creates worlds of truths to serve him as refuges where he can hide from the responsibility of the freedom that is his precisely because he inhabits the sensible, empirical world and *not* a world of fixed truths (where he would, clearly, be *less* free).

45. Sartre (1946), p. 16. Translation mine.

CONCLUSION

This conception of things seems to find confirmation in the classical experiments conducted into cognitive dissonance: the cognitive agent caught in such a situation creates for herself another world in order to attenuate the tension of incoherency she is experiencing in the empirical world. To be authentic, however, is to knowingly and wilfully confront this tension head on, to accept its consequences, to transform the situation into a step towards the realisation (the *making real*) of one's personal project, a project which encompasses one's entire *living*: "Man is nothing but his project, he exists only insofar as he realises himself, he is therefore nothing other than the full set of his acts, nothing other than his life."⁴⁶ From here, we can affirm that to live one's life as an authentic project means precisely the same thing as living one's life as a *telos*. In the end, it should not surprise us to discover that Sartrean authenticity and the Stoic *telos*, as we saw it redefined in the first section, have fused together over the course of this study. Simply put, one provides an answer to the Gettier problem the other to the resolution of cognitive dissonance problem, where both of these, we now see, are just two sides of the same epistemic, but above all metaphysical, coin: man's effort to assimilate his environment through the universal creation of virtual, mental worlds which are coherent abstractions of the empirical world. If the hypothesis be accepted, then the challenge, as I have tried to outline it, is to find ways of using this universal, demiurgic capacity in the most authentic way possible.



46. Sartre (1945), p. 51. Translation mine.

Bibliography

ANDERSON, Craig A., LEPPER, Mark R. & ROSS, Lee
— 'Perseverance of Social Theories: The Role of Explanation in the Persistence of Discredited Information', *Journal of Personality and Social Psychology*, Vol. 39, No. 6, 1980, pp. 1037-1049.

AYER, Alfred J.
— *The Problem of Knowledge*, London: Macmillan, 1956.

CHISHOLM, Roderick
— *Perceiving: A Philosophical Study*, Ithaca, NY: Cornell University Press, 1957.
— *Theory of Knowledge*, Englewood Cliffs, NJ: Prentice Hall, 1966.

CICERO
— *De Finibus*, quoted and translated from the bilingual (Latin-French) version available on www.remacle.org (still online in February 2017).

DUTANT, Julien
— 'Pourquoi le problème de Gettier est-il si important?', *Klesis*, No. 9, 2008, pp. 63-102.

FESTINGER, Leon
— 'Social Psychology and Group Processes', *Annual Review of Psychology*, 6, 1955, pp. 187-216.
— *A Theory of Cognitive Dissonance*, Stanford, CA: Stanford University Press, 1957.
— 'Cognitive Consequences of Forced Compliance', *Journal of Abnormal and Social Psychology*, Vol. 58, No. 2, 1959, pp. 203-210.

FESTINGER, Leon, RIECKEN, Henry & SCHACHTER, Stanley
— *When Prophecy Fails*, Minneapolis, MN: University of Minnesota Press, 1956.

GETTIER, Edmund
— 'Is Justified True Belief Knowledge?', *Analysis*, No. 23, 1963, pp. 121-123.

GOLDMAN, Alvin
— 'A Causal Theory of Knowing', *The Journal of Philosophy*, No. 64, 1967, pp. 357-372.

BIBLIOGRAPHY

— 'What Is Justified Belief?', *Justification and Knowledge*, Pappas (ed.), Dordrecht, The Netherlands: Reidel, 1979, pp. 1-23.

— *Knowledge In A Social World*, Oxford: Oxford University Press, 1999.

GOODMAN, Nelson

— *Ways of Worldmaking*, Indiana, IN: Hackett Publishing, 1978.

GRECO, John

— 'Virtues and Vices of Virtue Epistemology', *Canadian Journal of Philosophy*, No. 23, 1993, pp. 413-432.

— 'Knowledge as Credit for True Belief', in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, Depaul & Zagzebski (eds.), Oxford: Oxford University Press, 2003, pp. 111-134.

HARMAN, Gilbert

— *Change In View: Principles of Reasoning*, Cambridge, MA: MIT Press, 1986.

HEIDER, Fritz

— *The Psychology of Interpersonal Relations*, New York, NY: Wiley, 1958.

HETHERINGTON, Stephen

— *Knowledge Puzzles: An Introduction to Epistemology*, Boulder, CO: Westview Press, 1996.

— 'Actually Knowing', *Philosophical Quarterly*, No. 48, 1998, pp. 453-469.

— *Good Knowledge, Bad Knowledge: On Two Dogmas of Epistemology*, Oxford: Oxford University Press, 2001.

— 'Knowledge That Works: A Tale of Two Conceptual Models', in *Aspects of Knowing*, Hetherington (ed.), Amsterdam: Elsevier, 2006, pp. 219-240.

— 'Knowing-That, Knowing-How, and Knowing Philosophically', *Grazer Philosophische Studien*, No. 77, 2008, pp. 307-324.

HUME, David

— *An Enquiry Concerning Human Understanding* (1748), Oxford: Clarendon, 1975.

KLEIN, Peter

— 'A Proposed Definition of Propositional Knowledge', *The Journal of Philosophy*, No. 68, 1971, pp. 471-482.

BIBLIOGRAPHY

LA FONTAINE, Jean de

— 'Le Renard et les Raisins', *Fables*, Book III, in *Œuvres complètes*, Vol. I, Paris: Gallimard (Bibliothèque de la Pléiade), 1991, p. 124.

LEHRER, Keith

— *Theory of Knowledge*, Boulder, CO: Westview, 1990.

LEHRER, Keith & PAXSON, Thomas

— 'Knowledge: Undefeated Justified True Belief', *The Journal of Philosophy*, No. 66, 1969, pp. 225-237.

PLATO

— *Apology*, in *Plato: Complete Works*, Cooper & Hutchinson (eds.), Indianapolis, IN: Hackett, 1997, pp. 17-36.

— *Meno*, in *Plato: Complete Works*, Cooper & Hutchinson (eds.), Indianapolis, IN: Hackett, 1997, pp. 870-897.

— *Phaedo*, in *Plato: Complete Works*, Cooper & Hutchinson (eds.), Indianapolis, IN: Hackett, 1997, pp. 49-100.

— *Theaetetus*, in *Plato: Complete Works*, Cooper & Hutchinson (eds.), Indianapolis, IN: Hackett, 1997, pp. 157-234.

RODIER, Georges

— *Études de philosophie grecque*, Paris: Vrin, 1957.

ROSS, Lee & ANDERSON, Craig

— 'Shortcomings in the Attribution Process: On the Origins and Maintenance of Erroneous Social Assessments', *Judgement under Certainty: Heuristics and Biases*, Cambridge: Cambridge University Press, 1982, pp. 150-151.

RUSSELL, Bertrand

— *Human Knowledge: Its Scope and Limits*, London: Routledge, 1948.

RYLE, Gilbert

— 'Knowing How and Knowing That', *Proceedings of the Aristotelian Society*, No. 46, 1945, pp. 1-16.

— *The Concept of Mind*, Chicago, IL: University of Chicago Press, 1949.

SARTRE, Jean-Paul

— *Esquisse d'une théorie des émotions* (1938), Paris: Hermann, 1960.

— *L'Être et le Néant. Essai d'ontologie phénoménologique* (1943), Paris: Gallimard (Tel), 1994.

BIBLIOGRAPHY

— *L'Existentialisme est un humanisme* (1945), Paris: Gallimard (Folio), 1996.

— 'La liberté cartésienne', in *Descartes*, Genève: Traits-Trois Collines (Les classiques de la liberté), 1946. (The same text also appears in *Situations 1*, Paris: Gallimard, 1947.)

SOSA, Ernest

— 'The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge', *Midwest Studies in Philosophy*, No. 5, 1980, pp. 3-26.

— 'The Coherence of Virtue and the Virtue of Coherence: Justification in Epistemology', *Synthese*, No. 64, 1985, pp. 3-28.

— *Knowledge in Perspective: Selected Essays in Epistemology*, Cambridge: Cambridge Studies in Philosophy, 1991.

— *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Oxford: Oxford University Press, 2007.

SOSA, Ernest & TURRI, John

— 'Virtue Epistemology', in *The Encyclopedia of Philosophy and the Social Sciences*, Los Angeles, CA: Sage, 2013.

THÉBERT, Angélique

— 'La connaissance comme capacité: Le traitement des cas de Gettier', *Igitur*, No. 4, 2012, pp. 1-25.

TURRI, John

— 'Manifest Failure: The Gettier Problem Solved', *Philosophers' Imprint*, Vol. 11, No. 8, 2011.

— 'Is Knowledge Justified True Belief?', *Synthese*, Vol. 184, No. 3, 2012, pp. 247-259.

VIALA, A., OTA, H., VACHERON, M.-N., MARTIN, P. & CAROLI, F.

— 'Les Japonais en voyage pathologique à Paris: un modèle original de prise en charge transculturelle', *Nervure*, Vol. 17, No. 5 (suppl.), 2004, pp. 31-34.

ZAGZEBSKI, Linda

— *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, Cambridge: Cambridge University Press, 1996.

— 'Responses', *Philosophy and Phenomenological Research*, Vol. 60, No. 1, 2000, pp. 207-219.

Correspondence
metaphysiology@gmail.com

Cover illustration: Amazon archer drawing an arrow from her quiver as she turns to shoot at the enemy, as depicted on an Attic plate painted by Epiktetos, *ca.* 520-510 BCE (British Museum). Page layout and cover design of this digital book by Clément Bresch, February 2017.