KNOWLEDGE AS SUCCESS FROM ABILITY

By

Marina Bakalova

Submitted to
Central European University
Departament of Philosophy

In partial fulfilment of the requirements for the degree of Doctor of Philosophy
in Philosophy

Supervisor: Professor Katalin Farkas

Budapest, Hungary
2010

ABSTRACT

This dissertation attempts to offer an analysis of knowledge. The investigation is centered on a minimal sufficient condition for knowledge, a condition that is also supposed to account for knowledge in its full generality at a basic level (not including levels of epistemic ascent). My starting points are the theories of virtue epistemology developed by Ernest Sosa and John Greco. I defend the view that knowledge is success from intelligent ability (the KSA thesis), including both practical and theoretical abilities. I argue that KSA is a sustainable definition of knowledge. The following elements in my analysis of knowledge are crucial: knowledge is a broad mental state; knowledge is analyzable into internal and external components; knowledge entails credit; knowledge is compatible with various forms of epistemic luck. I consider each one of these points separately. The last chapter examines the application of the view to the problem of perception.

AKNOWLEDGMENTS

I would like to thank my supervisor Katalin Farkas for her acute comments and useful advice on the work in progress. Also, I would like to thank Nenad Miscevic for the discussions on the topic throughout the years, and John Greco for his directions and help with my work during my Doctoral Research Support Program at Fordham University in New York. I would like to express my gratitude to the faculty members of the Philosophy Department at CEU for their efforts to raise my professional level during these years.

TABLE OF CONTENTS

INTRODUCTION

There is a general pessimism in contemporary philosophical literature concerning the possibility of analysis of knowledge. This pessimism is due partly to the history of failures in attempting to define knowledge, and partly to a powerful argument offered by Timothy Williamson (2000, pp 2-5, 27-33 and 65-89) against the possibility of analyzing knowledge. We have at least two *prima facie* reasons to try to resist this pessimism, if possible. First, definitions are natural targets of our intuitive *a priori* philosophical thinking, and by being such, they direct one's thought to what is most general and universal, and at the same time to what is immune to counterexamples. These are rare achievements in philosophy, and yet desirable aims of philosophical thought. Second, introducing irreducible primitives into a philosophical theory is always problematic; some would say a sign of desperation. So, it seems that reconsidering the possibility of providing an analysis of knowledge is worthwhile. In what follows, I try to vindicate, and elaborate upon an already existing analysis which takes knowledge to be success from ability. The analysis is inspired by Ernest Sosa's definition of knowledge as apt belief. A belief aptly formed is a belief which is true *because* competent. If one is to know, one has to be a competent cognizer, and one has to reach the truth *as a result* of one's competence.

Why do I take this particular analysis as a starting point? The definition of knowledge as apt belief has three main virtues. First, it combines a rather minimal notion of credit with direct access to the objects of knowledge. Most of the research in epistemology has been about non-accidentality conditions for knowledge, centered upon the question what makes a belief non-accidentally true. I think that the most intuitive non-accidentality condition lies in a really minimal notion of credit, close to the one offered by virtue epistemologists, but arguably not identical to it. I think that epistemic credit is due to the refinement of one's cognitive system in such a way which provides one with a direct access to the objects of

knowledge. Second, Sosa's definition is wide enough to encompass all instances of knowledge – not just theoretical, but also practical. Thus, the definition can satisfy those who think that practical knowledge is not reducible to theoretical knowledge, and they seek for an account of practical knowledge too. Last, but not least, the definition is immune to the main counterexamples against analysis of knowledge in contemporary literature, or so I shall try to argue.

What is minimally sufficient for knowledge is a central question that this dissertation tries to answer. After the failure of the classical definition of knowledge as justified true belief due to the Gettier cases, generic reliabilism came to the scene as an alternative account of knowledge. According to generic reliabilism, knowledge is true belief which results from a *de facto* reliable process. The *de facto* reliability requirement, however, has been proven to be too weak for knowledge, as revealed by a series of counterexamples such as the Thermometer (David Armstrong 1973), Mr. Truetemp (Lehrer 1990), Clairvoyant case (BonJour, 1985), etc. In the recent decades, virtue epistemology has been developed partly as an attempt to improve upon generic reliabilism in various respects. One such respect is an attempt to rehabilitate a more sophisticated form of reliabilism by building in a subjective element into the *de facto* reliability requirement. The subjective element comes in the form of 'intellectual virtue'. Knowledge is claimed to result from the exercise of intellectual virtue instead of just a reliable process. Intellectual virtue is something for which we credit the agent, and so an agent's credit in achieving truth is set up as a crucial requirement for knowledge. Depending on the theory, the notion of credit varies from mere attributability (Sosa 2007) to the strongest requirement of responsibly sustained intellectual character (Zagzebski 1996, pp 59-69, Code 1984). Also, virtue can be seen as bearing different relations to *de facto* reliability. For instance, it can be seen as being built in the reliability requirement itself (a proper function of our cognitive faculties), or a bit more than that – cognitive integration of our faculties (Breyer

and Greco 2008), or as being independent of reliability (Montmarquet 1993), or as thick requirement over and above reliability in terms of possession of responsibly sustained intellectual character (Zagzebski 165-197). I claim that the notion of credit is an essential anti-accidentality condition. For the purpose of defining minimally necessary and sufficient conditions for knowledge, virtue is more than what we need, though. Throughout the text, I present four examples of knowledge without virtue. These include 1. knowledge from a weak module; 2. knowledge from improving ability; 3. knowledge obtained unwillingly, or even contrary to one's will; and 4. knowledge obtained through "savant syndrome". I suggest that a minimally necessary justifier of our beliefs is an intelligent act, and a minimally necessary condition for knowledge is an intelligent ability developed through a sequence of intelligent acts. The notion of intelligent ability involved in my account could be seen as close to the Platonic notion of virtue as proper function, but it is arguably different. For one thing, I claim that there is no need to refer to virtues over and above mere intelligent abilities.

What is an intelligent ability? I suggest an understanding of intelligent ability close to Gilbert Ryle's understanding of knowledge how (Ryle 1949, pp 40-50). Intelligent ability is an ability refined through experience. It is an ability which results from one's conscious or sub-personal investment of propensities, i.e. paying heed to certain factors relevant to one's cognitive task. It is not a physiological ability like the ability to breathe or digest, or like a thermometer's ability to measure the temperature. Hence, we can see why the definition of knowledge as success from intelligent ability can apply to practical knowledge too. Both knowledge how and propositional knowledge are based on exercise of intelligent abilities of the purported kind.

Is the agent responsible for her knowledge? Together with virtue epistemologists, I agree that knowledge is an achievement. However, it is not necessarily an achievement creditable to the agent, but it is always an achievement creditable to her cognitive system. For

instance, we have intelligent abilities which we did not want to have. Some of us did not want to go to school, and simply by being forced to do it they learned how to count, read, etc. Propensities that essentially drive intelligent acts do not amount to virtues, since they could be invested into acquiring intelligent abilities without the agent being motivated to acquire them. Hence, a part of our intelligent abilities is not an achievement for which we are to credit the agent as a conscious being. An example of this sort is perception. During perceptual learning the agent's eyes explore the environment, and thanks to that process of exploration the agent obtain the ability to visually discriminate, and recognize objects directly. But this exploration happens mostly without the agent's awareness. One's cognitive system learns how to perceive, but it is not the agent who controls the process. So, I take it that the minimal competence entails an intelligent ability sub-personally acquired.

Another important point in my dissertation is to try to estimate the role of epistemic luck involved in achieving success through ability. This relates to the following question: in what circumstances do we know? First, in the process of acquiring an intelligent ability, that ability is being attuned to a set of circumstances. I call these "normalized circumstances" for the ability. Their presence is necessary for knowledge. Normalized circumstances involve some risks of failure. I claim that these risks can be graded, and they give rise to various forms of luck. Most of these forms of luck do not affect one's knowledge by simply being present in one's cognitive situation. They affect knowledge only when they yield to error.

Let me now provide a brief overview of the content of the separate chapters.

In chapter 1, I introduce the reader into the background of the thesis, and I settle some preliminaries. I explain the motivation for Ernest Sosa's virtue epistemology. Also, I discuss the two main lines in virtue epistemology. At the end, I sketch a general argument to prepare

the ground for the rest of the thesis. There, I argue that virtues are more than what we need to define knowledge.

In chapter 2, I explore the potential of providing a metaphysical analysis of knowledge based on the notion of success from ability. My aim here is to extend Timothy Williamson's defense of primeness of knowledge (Williamson 2002, pp 2-5, and pp. 49-93) to practical abilities, and to show that the notion of knowledge as success from ability is a prime notion. I defend a view, according to which knowledge is accepted to be a broad mental state, and at the same time it is claimed to be analyzable.

Chapter 3 is a defense of the claim that justification is not separate from skills and capacities, and that skills and capacities constitute the essential basis, that is both temporarily and normatively prior to the traditionally appreciated knowledge-that. I suggest that Gilbert Ryle provided a rather good theory of what intelligent ability is in his account of knowledge how. He was right to think that the exercise of these abilities is of fundamental significance, and that they underlie propositional knowledge too. In the same chapter, I suggest that Ryle foresaw some of the basic ideas associated with the contemporary virtue epistemology.

In chapter 4, I offer a theory of epistemic luck which sets out in detail which kinds of luck are compatible with knowledge and which ones are not. I divide the issue in two parts: luck compatible/incompatible with the possession of the internal basis of ability, and luck compatible/incompatible with the presence of appropriate circumstances. In this chapter, I argue against Duncan Pritchard's anti-luck epistemology and in favor of the sufficiency of the thesis that knowledge is an achievement. I claim that the crucial anti-luck condition is already integrated into the achievement thesis. Other kinds of luck are compatible with knowledge.

In chapter 5, I explore the application of the theory defended throughout the thesis to the problem of perception. I suggest that my arguments provide a further support for Alan Millar's theory about the role of perceptual recognitional abilities in perceptual knowledge

(Millar 2008 a, 2008 b, 2010 in Pritchard, Millar, Haddock). I try to defend a version of ability disjunctivism about perceptual knowledge based on narrow competence.

CHAPTER 1: VIRUE EPISTEMOLOGY

In this chapter, I introduce the reader into the background of the thesis, and try to settle some preliminaries. The chapter consists of five sections. In section 1, I outline very briefly the main idea of virtue epistemology. In section 2, I place virtue epistemology in the context of traditional and contemporary epistemological debates, like those between internalism and externalism. In section 3, I present Sosa's Two-Level Epistemology. In section 4, I discuss the two main lines in VE, and I try to explain in what sense we can talk about virtue when we refer to mere faculties. In section 5, I sketch a general argument to prepare the ground for the rest of the thesis. I suggest that we take possession of intelligent skills rather than possession of intellectual virtues as a minimal necessary condition for knowledge.

1.1 What is virtue epistemology?

Virtue epistemology starts from the widely accepted assumption that epistemology is a normative discipline, although it does not neglect its descriptive underpinnings. Normative epistemology, as Jaegwon Kim (1988, p. 381) puts it, is centered upon the question of what people *ought* to believe. These questions concern what the norms of justification are; they entail formulating recommendations for generating and revising beliefs, answering to the skeptic, etc. The assumption that epistemology is a normative discipline can be expressed as an endorsement of the role of epistemology in answering these questions.

A reason why such an endorsement is needed is mainly due to Quine's famous paper "Epistemology Naturalized" (Quine 1969), where he attempts to undermine the enterprise of normative epistemology. Quine is particularly skeptical about the possibility of providing a normative explanation of how our sensations (empirical data) justify our beliefs about facts.

Hence he concludes that epistemologists should abandon their proper subject - questions about what is reasonable to believe - and should become a chapter of cognitive psychology. His dire diagnosis is expressed by an often quoted passage:

> The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology? (Quine, 1969, p. 75)

This is a consequence that no epistemologist should be particularly happy to accept. Virtue epistemologists offer a way out of Quine's methodological project. The main idea of virtue epistemology is to provide an account of knowledge and justification by referring to the kind of normativity involved in virtue theories. The focus is on understanding the nature of epistemic norms, value and evaluation, and this focus is a defining feature of the field. Thus virtue epistemologists are mainly responsible for bringing forward for the recent "value turn" in epistemology (see Riggs 2006, Pritchard 2007). The way in which virtue epistemologists oppose the naturalized project is by disagreeing with the claim that traditional epistemology should be replaced with the psychological study of how we reason, or that epistemological statements could be explained exhaustively in terms of scientifically respectable objects and properties. However, some virtue epistemologists endorse more modest naturalized methods, according to which philosophers are allowed to make use of results from sciences studying cognition, and from history in order to resolve epistemological issues. (Greco 2001; Part II; Sosa 1991, 105–6; Zagzebski 1996, 336–7)

Sometime after the revival of virtue ethics during the 70s, Ernest Sosa appealed to the idea of intellectual virtues in his seminal works: "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge,"(Sosa, 1980) and in *Knowledge in Perspective* (Sosa 1991). Sosa drew attention to the intelligent dispositions of the agent as basis of achieving truth and as a primary source of epistemic value. Thus, he suggested a shift of the

epistemological focus. In contrast to previous theories focused primarily on beliefs, and their justificatory status, Sosa argued that:

> Primary justification would apply to intellectual virtues, to stable dispositions for belief acquisition, through their greater contribution toward getting us to truth. Secondary justification would then attach to particular beliefs in virtue of their source in intellectual virtues or other such justified dispositions. (Sosa 1980, p. 23)

Intellectual virtues set the norms of how the agent should form her beliefs in a certain set of circumstances. The proposal is that agent and her virtuous character should become the primary focus of epistemic evaluation.

Virtue epistemology has its roots in the Ancient Greek theories of intellectual virtues. The classical standard theory of intellectual virtues was offered by Aristotle (in his *Nichomachean Ethics*, book B (II) and above all Z (VI). He introduces epistemic (intellectual) virtues in the very beginning of his book B:

> Virtue, then, being of two kinds, intellectual and moral, intellectual virtue in the main owes both its birth and its growth to teaching (for which reason it requires experience and time), while moral virtue comes about as a result of habit, whence also its name (*ethike*) is one that is formed by a slight variation from the word *ethos* (habit).[1]

Moral and intellectual virtues are essentially distinguished by their goal, namely moral virtues aim at the good, whereas intellectual virtues aim at truth. The function of all virtues is to guide oneself to happiness. The virtues in the two domains are connected in this respect, but it is unclear exactly to what extent. Aristotle thought that the intellect is supposed to shape one's moral character; it must guide one's actions by suggesting the right motive, or the right desire. The intellectual virtues, on his theory, divide into two groups: contemplative ones, which include wisdom (*sophia*), intuitive reason (*nous*) and scientific knowledge (*episteme)*; and practical ones, which include practical reason, or *phronesis*, and craft (*techne*). The two parts of the intellect are meant to achieve truth, but it is practical reason, in particular, that

[1] The italic is mine.

guides the actions in accordance with the right desire. It is the one by means of which moral virtues are a combination of intellect and character. Contemplative virtues are taken to be independent of practical reason, and hence of the moral domain, because they can choose independently of desire. Speculative reasoning is about necessary truths, and contingent desire seems not to disrupt reasoning about what is necessary.

However, Aristotle approves that intellectual virtues are not just natural capacities, but they result from teaching in a similar way that moral virtues result from habituation, and therefore they are, like moral virtues, refined through practice. Linda Zagzebski thinks that Aristotle in vain tried to sustain the relative independence between intellectual and moral domains (Zagzebski 1996, pp. 211-231). Although contemplative virtues are about necessary truth, and do not involve desire, she argues, since Aristotle sets necessary truths as the end of contemplative reason, he must have recognized that without desire there is no reasoning (pp. 215-216). So, she argues that *phronesis* has a role to play in the contemplative domain too; it has to lay the right desire for truth. As a result of similar considerations some contemporary virtue epistemologists embrace Aristotle's account of the moral virtues as the model for understanding intellectual virtue. They take intellectual virtues to be traits of intellectual character like moral virtues are traits of moral character.

However, we encounter another use of 'virtue' in Aristotle's *Nicomachean Ethics*. On this use, virtue is a proper function of something. "The virtue of a thing is relative to its proper work." (Book VI, 2). Then, maybe, we can see contemplative virtues as being virtues in this sense: they could be just properly functioning cognitive faculties, and not character traits as moral virtues are. This understanding of virtue, which is more prominent in Plato, is adopted and applied to intellectual virtues by another group of contemporary virtue epistemologists: Ernest Sosa, and John Greco.

Let me return to the issue how contemporary virtue epistemology came on the scene, and explain why Sosa referred to intellectual virtues. As far as I understand, he has three general reasons to involve the notion of intellectual virtue in the epistemic enterprise. First, by doing so, Sosa wants to introduce a normative element into the 'simple' reliabilist theory of knowledge, and thus to turn it into a theory with its own normative apparatus, yet open to scientific discoveries concerning human cognitive processes. Second, Sosa have seen a potential in the notion of intellectual virtue to enable a project of unifying between various camps in epistemology, including a kind of foundationalism, against Quine's above-mentioned complaint. Finally, a virtue-centered epistemology brings insights into solving some long-standing problems in epistemology, such as the value problem, the problem of skepticism, the more recent Gettier problem, etc.

I will present Sosa's basic idea of virtue epistemology in the next two sections. First I will try to explain the motivation of Sosa's unification project, which is central to his epistemology, and then I will focus on presenting Sosa's view itself.

1.2 Placing Sosa-style virtue epistemology on the epistemological map

The project of unifying internalism and externalism in epistemology is central to Sosa's epistemology. In this section, I will try to explain the rationale for such a project. I will proceed by drawing a map of traditional epistemological positions, and present the general shortcomings of each one of them taken separately, involving some of Sosa's own criticisms. The exposition should clarify why they need to be unified.

1.2.1 Internalism vs. externalism

The most general distinction between epistemological views is the one between internalism and externalism. Here is a brief reminder of the distinction. Most, if not all

theories of knowledge take for granted that knowledge is something more than mere true belief, thus their main task is to provide and defend a detailed account of the condition, the satisfaction of which turns merely true beliefs into knowledge.[2]

On one approach, the epistemic condition is construed in a way that the agent must have some sort of reflective access to the factors that render her belief justified. This view is called epistemic internalism. Its strongest forms bear a relation to the so called "KK principle" (H. A. Pritchard 1950) according to which in order to know, one is required also to be in a position to know that one knows. There are many varieties of epistemic internalism ranging from demands of stronger to those of weaker reflective awareness.[3]

The classification in terms of strength of the internalist requirement ranges along three dimensions. First, it varies depending on whether we demand actual or potential access to the justifying factors. Second, it depends on whether we require that the thinker has access to the fact that *J* in particular is a justifier of *p* or that she simply has (a potential) access to a justifier without knowing that it is the justifier for *p*. Third, it depends on whether we require access to all of the justifiers or only to some essential subset of them.

Thus, the strongest version of epistemic internalism will concern the claim that an agent stands in a positive epistemic relation to *p* only if she is actually aware of all the justifiers of *p*. For instance, Mary is justified in believing that Abraham Ibn Daud was a Spanish Jewish philosopher, because she in fact remembers all the evidence that she ever had in support of this proposition, e.g. she read it in the Cambridge Dictionary of Philosophy five minutes ago. Certainly, the range of propositions we are justified to believe in this way is very limited. Hence, in order to render the condition more comprehensive, internalists are pressed to make it more permissive.

[2] Here I provisionally leave aside the details of Timothy Williamson's account, which does not factorize knowledge in the same way. See Williamson (2000).
[3] For further references see George Pappas (2008).

The weakest version of epistemic internalism is the view according to which an agent stands in a positive epistemic relation to *p iff* she can potentially gain access to some essential justifier *J* of *p* without particularly knowing that *J* is a justifier of *p*. This refers to a case in which Mary remembers that Abraham Ibn Daud is a Spanish Jewish philosopher, but she does not remember how she learnt it. However, there is a possibility for her to check it out in the dictionary. There is a range of views between the strongest and the weakest forms of epistemic internalism, but we don't need to go in such detail.

Epistemic externalism, on the other hand, is the rejection of epistemic internalism. It is the view that the epistemic condition must be *de facto* instantiated in the production of an agent's belief, with no need of the agent's reflective awareness of the justifying factors. The most prominent example of epistemic externalism is reliabilism, the view that *S* is justified in believing *p iff* her belief that *p* is produced by a reliable process, where reliability entails truth in some substantial proportion of cases.

Now, depending on the particular view, the dividing line between internalism and externalism will shift. As Bonjour (1993, p. 133) points out, a weaker version of internalism would be externalist in relation to a stronger version of internalism. For example, the potential-access internalism would count as externalism in relation to the actual access internalism. So, as we can see, the dividing line between the two has become rather vague in the course of the ongoing debates. Here is one way to make sense of the distinction. The externalist's point is that *de facto* reliability of a process/faculty fulfils the epistemic condition, no matter whether the agent has potential or for that matter actual access to that condition or not. The view which takes both de facto reliability and access to the justifiers as necessary for justification will be a hybrid view. Potential access internalism differs from externalism in its major claim that potential access to the justifiers gives us justification independently of the fact whether *p* is reliably produced or not.

1.2.2 Foundationalism and coherentism

Now, let me turn to the distinction that originated within the internalist camp, the one between foundationalism and coherentism. Both foundationalism and coherentism are traditionally internalist theories of justification, although each of them can have its externalist counterpart (Sosa in BonJour and Sosa 2003, p.7). [4]

According to foundationalism, our body of knowledge consists of two groups of beliefs: basic beliefs (that are non-inferentially justified), and non-basic beliefs (that are justified by inference from basic beliefs). According to classical foundationalism, basic beliefs pertain to immediately experienced mental states, and all other beliefs stand in a cognitively accessible inferential relation to those beliefs. Depending on the particular foundationalist story, basic beliefs are either self-justified or justified by being based on experience, i.e. by our basic awareness of one's experiential states. Accessibility to the justifying components is taken to be crucial for both inferential and non-inferential justification, and for this reason, classical foundationalism is considered to be an internalist theory of justification.

The epistemic status of basic beliefs is famously a problematic part of the foundationalist theories. To clarify the status, we need to answer the question of how non-inferential justification is possible. Non-inferential justification is usually taken for granted and it refers to some sort of direct relation between a belief and its truth, which is not mediated through any other belief. This relation is usually explained in terms of privileged *access* which one has to one's own mental states and which one does not have to others' mental states. My privileged access to my own mental states guarantees the immediate

[4] Externalist foundationalism would be one with reliabilist foundation and *de facto* inferential links to the basic beliefs of which the agent need not be aware. An externalist version of coherentism would be the view that justification draws from a *de facto* coherent relations between the justificandum beliefs and the justifying beliefs.

relation to the truth of my beliefs about my mental states. For instance, my belief that I am in pain is so intimately tied to my being in pain that I can hardly go wrong about it. The same holds for my belief that it seems to me that there is a forest out there (which is a belief about my own mental state, uncommitted to the truth of there being a forest out there).

A famous *reductio* argument against the possibility of basic beliefs is raised by Bonjour (1985) who claims that putatively basic beliefs are justified only if we have a reason to believe that they are basic beliefs, and that beliefs of that sort are justified. So, justification of a putatively basic belief depends on further beliefs. Therefore, basic beliefs are impossible.

Another problem related to basic beliefs is that their content is too narrow to support the whole body of knowledge. The content of basic beliefs is usually limited to seemings and looks. The foundationalist cannot simply assume that these seemings provide us with a straightforward access to the external facts because of the well known skeptical challenge: I don't know that what seems to me is really out there until I rule out the possibility that I am cleverly deceived. But if I am to rule out this possibility even in general and for once, I have to rely again on appearances, and the basic beliefs thereof, and it all becomes circular. In the face of this challenge, basic beliefs give us at best very scarce information about the external world, and cannot support our beliefs about the external world. But these beliefs constitute the major part of our knowledge. Therefore, it remains unclear how basic beliefs can serve as a foundation rich enough to support the whole body of our knowledge. Their content is simply too meagre, and their power to justify is not sufficient to play this role.

Greco and Turri (2010) suggest a further dilemma concerning the foundationalist's epistemic principle. The epistemic principle about how perceptual experience justifies a perceptual belief can be understood in two different ways: either as a fundamental principle about epistemic justification, or as an instance of a more general principle. If it is understood

in the former way, the principle disintegrates into a multitude of fundamental principles for each particular sensory modality. These principles would lack a unifying ground. If the principle is understood as a more general principle, encompassing all sensory modalities, the founationalist will have to account for a deeper unifying ground.

Let me now turn to classical coherentism. Coherentism can be a theory of truth, a theory of justification, or a theory of the content of beliefs. We are interested in coherentism about justification only. Coherentism about justification is traditionally an internalist view, since it requires from the believer to have access to the coherence relation of her belief to other beliefs in her system of beliefs. According to coherentism, one's belief is justified *solely* by its coherence with the rest of the body of beliefs.

A major problem with coherentism is that a perfectly coherent web of beliefs can be isolated from reality. A good example of that is the victim of Descartes evil demon, whose beliefs about the purported reality are coherent, but none of them amounts to knowledge. A related problem is that our perceptual beliefs, for instance, have few connections with other beliefs in the total system. If I now believe truly that my cat is sleeping on the sofa, and I replace that belief with its negation, this would have very little effect on the overall coherence of my system of beliefs. Indeed, my new belief can be exactly as coherent with the rest as the old one, and therefore, according to coherentism, it would be equally well justified. We do not want our theory of justification to have such consequences. So, the basic hurdle for the coherentist view is the connection between internal relations (justification) and external targets and properties of belief, above all its truth.

A second problem concerning the positive coherentist framework is how one justifies the very first belief that one ever forms. On the more practical side, it seems that if one is to follow the coherentist norm of justification strictly, one must suspend judgment for quite a while until one obtains a rich enough web of beliefs. A subsequent theoretical problem is that

the coherentist falls into a vicious circle, when questioned how coherence ever enters into the system of beliefs. She must say that it enters into the web of beliefs when there are sufficiently many coherent beliefs.

The alternative is weak coherentism but it is already, as I said, a hybrid view. If we accept a background theory about our relationship to the world, we need to trust our perception of facts about the world on the first place. But in order to trust it, it has to have some trustworthy quality independent of coherence, such as reliability.

Let us now turn to externalism, the view according to which the necessary and sufficient condition for justification of our beliefs is that they are reliably formed. The main externalist theory is generic reliabilism. This is the theory that a belief *p* is justified if it results from a *de facto* reliable cognitive process. A *de facto* reliable process is a process which produces truth in most of the cases of belief-formation. John Greco (2002, p. 291) points out two main advantages of generic reliabilism. First, it explains why beliefs based on guessing cannot amount to knowledge. Second, it provides a resource against the skeptical argument. The skeptic requires a proof of reliability of one's faculties or belief-forming processes: if one is to know that *p,* one must rule out the skeptical hypothesis, for instance, that one's senses are deceived by an evil demon. This means to prove that one's senses are reliable. Since such a proof cannot be established on the basis of internalist justification, the skeptical conclusion follows. Reliabilism requires just *de facto* reliability of one's belief forming processes in our actual world, and thus it cuts off the requirement of proving one's reliability that yields to the skeptical conclusion. (see Greco 2000)

However, generic reliabilism has been a subject of serious attacks. One problem of generic reliabilism is expressed by Bonjour's famous clairvoyant objection, which aims to show that reliability is not sufficient for knowledge (BonJour, 1985, p. 42). This objection concerns the story of Norman who possesses the reliable clairvoyant ability to tell where the

President of the United States is, but he has not the slightest inkling that he possesses this ability. One day, Norman suddenly finds himself believing accurately, on the basis of his clairvoyant ability, that the President is in New York (*p*). Although he has no evidence either way on whether *p* is true, on whether he possesses a clairvoyant ability, and on whether such ability is even possible, reliabilism commits us to holding that Norman knows that *p*. Yet, BonJour claims that Norman does not know that since, in light of his own subjective conception of the situation, believing that *p* is unwarranted. Therefore, it is accidental that his belief is true.

Another objection raised against reliabilism is supposed to show that reliability is not necessary for justification. According to this objection, it is plausible to think that the victim of Descartes evil demon is internally justified in what she beliefs on the basis of her sense impressions, despite the fact that none of her corresponding beliefs is reliably formed. The objection is known as "the new evil demon problem", and was first introduced in the literature by Cohen, S. and K. Lehrer (1983), and Cohen (1984). Later on, it was endorsed in an attempted solution by Comesana (2002) and Sosa (2003: 159-61).

Finally, a rather important problem raised against reliabilism is the so called 'bootstrapping problem'.[5] The bootstrapping problem is the problem that the reliabilist theory allows for someone who forms a reliable belief to know in a trivial way, i.e. by bootstrapping that her belief is reliable. But forming a belief in a reliable way, and knowing that one's belief is reliable are two rather different things. The problem is that reliabilism has no resources to account for the latter kind of knowledge.

Having in mind all this, it seems that a theory that combines the advantages of internalism and externalism without their disadvantages is exactly the one that we need.

[5] For presentation of the problem see Vogel (2008).

1.3 Sosa's two-level epistemology. The dialectics between externalism and internalism, foundationalism and coherentism

All these classical and well-discussed difficulties besetting each of the discussed warring factions in epistemology prompt the thought that reconciliation is needed, and that a project of unifying them might have a chance to avoid the difficulties. Sosa's original conception of virtue epistemology is an attempt to reconcile and to find the right place of the warring factors in epistemology outlined above. I have in mind, in particular, Sosa's attempt to do so by distributing and uniting the warring factions between two layers of our epistemic assent. In what follows I present Sosa's two-level epistemology.

1.3.1 Animal knowledge

The first layer of knowledge, according to Sosa, is the level of cognition. On that level, we simply exercise our cognitive faculties: perception, memory, intuition (or introspection), reason, and induction. [6] These are reliable faculties, which Sosa calls "intellectual virtues". Sosa uses a Platonic notion of virtue, according to which everything with a function has a virtue if it performs its function well. Our faculties are in this sense virtues or powers to learn various propositions, and to retain knowledge. They are naturally endowed to a species and they form a non-deviant cognitive structure. Sosa invites us to assume that the functioning of our faculties or virtues is truth-conducive and reliable provided that there is no special reason for caution: abnormal circumstances, or specific cognitive defect peculiar to the cognizer. Sosa further assumes that in most ordinary circumstances there is indeed no such reason. So, our reliable faculties are takes to provide the norm of justification on the first level.

[6] The theory is meant to apply to human beings. Other cognitive beings could possess different cognitive dispositions that would count as virtues in their epistemic communities.

The outcome of the functioning of our first-level virtues is some kind of rudimentary 'knowledge', which Sosa calls 'animal knowledge'. Let me explain what animal knowledge is. Animal knowledge is automatically produced by our well-functioning cognitive faculties; it requires an experiential state and some primitive conceptual apparatus. Supposedly, at least some animals have both, so the term 'animal' is not to be taken as entirely metaphorical. Here are a few examples of possible instances of animal knowledge: "The coffee is hot", "54+32=86", "I dropped my pen somewhere", "Fernando Trueba directed the movie "The Girl of Your Dreams". The specific characteristic of animal knowledge is that one animally knows that *p* without being in a position to explain on the basis of what one knows that *p*.

The overall animal beliefs based on our faculties or virtues can be seen as forming the foundation of knowledge. Thus Sosa offers a unifying ground to the foundationalist, based on the normative properties of intellectual virtue understood as reliable faculty. The notion of virtue plays a unifying role, because it has a general structure applicable to different modes of belief formation and different circumstances. Animal beliefs play the role of the foundationalist's basic beliefs. However, they are not self-justified, but justified on the basis of reliability of our faculties.

Sosa substitutes the foundationalist basic epistemic principle of privileged access with a reliabilist one not just to unify the foundation, but also to straighten the relation of our basic animal beliefs to external facts. The reliabilist principle gives advantage to the theory, because once the *de facto* reliability criterion is satisfied, the theory is freed from the skeptical burden to rule out the skeptical hypothesis. The relabilist criterion grants a relation to external facts in most cases.[7]

Another thing to notice is that basic and many non-basic beliefs from the classical foundationalist structure are part of animal knowledge, and thus distributed on the first level

[7] For further debate with the classical foundationalist, see Sosa's discussion with Richard Feldman in Greco (Ed.) 2000, pp 42-57, and pp. 287-90; and BonJour and Sosa 2003, pp 119-141.

of Sosa's two-level epistemology. Perceptual beliefs and beliefs of intuition form the group of basic beliefs, and beliefs of memory and reason are inferred or derived from the basic beliefs. All these beliefs, according to Sosa, are immediate productions of various faculties, and thus are placed on the first level. So, we have many beliefs from the upper level of the classical foundationalist pyramid included in animal knowledge.

1.3.2 Reflective knowledge

Clearly, human subjects are capable of deeper and more sophisticated knowledge than that which is spontaneously produced by their cognitive capacities. Take, for instance, one's reflective assurance that one is right in what one believes to be the case. This awareness seems to essentially involve a second-order perspective upon one's capacities. Sosa thinks that typical human knowledge involves such a perspective. Our beliefs produced on the first level are being justified by rational reflection upon the reliability of their sources. So, he proposes us to add one more level of knowledge on the top of 'animal knowledge' which he calls 'reflective knowledge'. Reflective knowledge introduces an internalist perspective of the cognizer and thus accommodates epistemic internalism, explained above.

Sosa suggests that the more characteristic type of justification for reflective knowledge is determined by the coherence of the overall body of beliefs, including beliefs about the reliability of one's particular cognitive faculties. Thus Sosa (1980) accommodates coherentism on the second level of the structure of knowledge. Coherence is the manifestation of intellectual virtue on that level; it determines the specific intellectual value of reflective knowledge.

Introducing coherent perspective has many purposes. The detachment of the levels in Sosa's epistemology, and the difference in criteria to which they are liable helps him to solve the bootstrapping problem for reliabilism. By keeping the epistemic criterion for animal

knowledge separate, Sosa achieves a system in which the method of reflection upon the criterion is not presupposed on the first level. Thanks to the enlightening reflection, Sosa claims, we gain more on the second level than we have on the first one. And vice versa, although the justifiedness of the first level knowledge is open to reflection on the second level, one does not need to be explicitly aware of the justifying factors on the first level, because there is a separate epistemic criterion for this sort of knowledge. (Sosa 1994) Thus, coherent reflection is supposed to guarantee an enlightening discovery over the reliability of one's sense organs, reason and other faculties, and not just a "blind" trust in them.

There are two related benefits. First, we obtain a virtuous circle of epistemic coherence in contrast to the circle in which the classical coherentist is trapped. The classical coherentist is bound to say that the rules of coherence involved in justification of beliefs are in turn defined as justified beliefs in terms of coherence, and thus she presupposes the principle, which she wants to define. On Sosa's theory, the basic epistemic criterion is the criterion of reliability, and our further epistemic principles have to satisfy this criterion. [8] Second, by complementing the coherentist web with bottom-up epistemic relations, Sosa solves the isolation problem for coherentism. The coherentist web is connected to the external world through the exercise of our reliable faculties.

Reflective perspective is supposed to play another important role too. It is offered in order to solve some problems of generic reliabilism, such as the generaily problem, and the clairvoyant objection. The generality problem is the problem that a token cognitive process is

[8] Van Cleve (1979) nicely expresses a similar idea by offering a three-layered structure of our epistemic assent. We first get to know non-epistemic propositions, call them p. Then we learn epistemic propositions, singular propositions that attribute evidence, certainty, etc. to the non-epistemic propositions, call them Ep. At the end, we come to believe epistemic principles, call them EP. For instance:

p: *a is F.*

Ep: *I believe that a is F, because I perceive it clearly and distinctly*

EP: *I am justified in believing all and only those things which I perceive clearly and distinctly.*

If we do not require normative justification for non-epistemic propositions like p, we end up with a virtuous circle of warrant from which we gain epistemic result, because our knowledge on the higher stage illuminates our initial beliefs.

an instance of several types, and it is unclear which type is relevant for evaluating reliability. Sosa argues that on the basis of inductive generalization from one's perspective, the relevant types are being traced out, although in a sketchy and implicit way. Concerning the clairvoyant example, Sosa claims that Norman does not know on the basis of his clairvoyant power only at an initial stage. According to Sosa, Norman needs to adopt a second-level perspective to gain the required certainty that his clairvoyant power is reliable, and to eliminate the apparent accidentality of his clairvoyant beliefs.

To summarize, the general picture of Sosa's two-level epistemology is that we have reliabilist (externalist) criterion for the first-level knowledge for our separate faculties, and a coherentist (internalist) criterion for the second-level knowledge. This is just the general picture. But indeed, the distribution of the jobs of the internalist and the externalist criteria is not so straightforwardly clear. Our internalist or reflective knowledge requires having access to the externalist justifier for our first-level beliefs. Partly, however, our reflective beliefs are also liable to the externalist criterion of reliability. The point seems more obvious in Sosa's recent work (Sosa 2008), where he introduced the aptness criterion of knowledge which applies equally to animal and reflective knowledge, and it is a kind of externalist criterion, which I discuss below in chapters 2-4. The point at which the externalist and the internalist criteria come together (and perhaps merge in a single criterion) is that when we require a cognizer to have access to her justifiers (from a coherent perspective), we require at the same time for this access to have a non-accidental internal basis which is reliable. So, the conclusion is that even though a cognizer can have animal knowledge without having access to her justifiers; having access to one's justifiers on reflective level requires the satisfaction of reliability criterion too. This might suggest that we shall opt for a unified epistemic criterion of reliability of our cognitive dispositions and leave the distinction between internalism and

externalism to philosophers of mind. This is one possible way to go, which I will pursue in the rest of the thesis.

Reliabilists like Goldman think that this epistemic-ascent proposal, especially raised in relation to the Clairvoyant case, sets too high a standard of justifiedness. Solving the Clairvoyant problem by involving an epistemic perspective suggests that first-level knowledge in general is dependent on the second-level perspective. We are bound to look at the small children in analogous way to Norman. It seems hardly plausible that small children gain subjective certainty for their beliefs by reflective assurance. For one thing, they do not have coherent reflective perspective, refined enough to provide them with such assurance. Young children have few if any such higher-order beliefs, but still have many first-order beliefs that are justified. Goldman's proposal (1986, pp. 111–112) [9] is to add a weaker supplementary requirement to *de facto* reliability in the form of a negative higher-order perspective, i.e. a non-undermining (or "anti-defeater") condition. It says that a cognizer must *not* have reason to believe that her first-order belief is *not* reliably caused.

Another critic of Sosa's virtue perspectivism is John Greco (2000a, pp. 187-190, and 2004, pp.96-106). He argues against the need to postulate a reflective perspective in order to solve problems related to first level-knowledge.[10] First of all, he thinks that to claim that people have Sosa-type reflective perspective is psychologically implausible:

> In typical cases of knowledge people do not have *any* beliefs about their beliefs, or beliefs about the sources of their beliefs, or beliefs about the reliability of those sources. (Greco 2004, p. 97)

Besides, Greco claims that one needs a highly detailed perspective to solve the generality problem and the meta-incoherence problem of the clairvoyant, which ordinary people lack. Therefore, it is something else that turns children's inborn capacities, as well as Norman's clairvoyant power into refined and trustworthy enough faculties. Greco argues that

virtue reliabilism has resources for coping with the problems which reflective perspective is meant to solve without appealing either to positive or to negative perspective. He proposes to bring the coherence criterion to the first level, and to investigate how it works on that level. In particular, the proposal is that we should focus on building a normative theory about the right kind of dispositions that provide the subject with sufficient insurance of the truth of his beliefs. His positive proposal concerns a theory of cognitive integration. I briefly explain the main idea of the theory in section 5 below.

I would like to note, however, that although Sosa's theory of reflective knowledge arguably does not provide the best solution of the clairvoyant examples; such knowledge is not of a stipulated kind. Our philosophical knowledge, for instance, is of that specific sort: it is an enlightening discovery over our first-level beliefs. Also, reflective knowledge sets a desirable height for the epistemic ascent of humans. So, I think it is still worth considering this kind of knowledge as a part of a metaphysical theory of knowledge. I shall refer to reflective knowledge in this sense in the latter text.

Despite the problems that Sosa's virtue perspectivism faces, it is a natural consequence of Sosa's unifying project that it substantially enlarges the means for solving various epistemological problems corrupting each view taken separately. My thesis is meant to reveal two major contributions of Sosa's theory that have not been given much attention. They are both related to his attempt to provide a metaphysical analysis of knowledge as apt belief. In chapter 2, I argue that Sosa's more recent aptness account of knowledge handles the grave consequences of Timothy Williamson's primeness argument against the possibility of analysis of knowledge. Next, in chapter 3, I argue that the same account gives us a potential to provide an account of knowledge in its full generality, covering not just propositional knowledge, but also practical knowledge.

---

[9] Michael Bergman (1997) defends the same view.
[10] For further criticism of epistemic perspectivism see Bernecker (2006), Lepock (2006).

1.4 The nature of intellectual virtues

So far, I have presented Sosa's general epistemology. We saw how the faculty view of knowledge combines the advantages of the previous views, and arguably avoids their weaknesses. Now, let me focus on explaining how contemporary theorists understand the nature of intellectual virtues and what their normative role in epistemology is. The most general and commonly accepted characterization of intellectual virtue is that it is a stable disposition to achieve truth in certain range of circumstances. On a strongest notion, virtue is a motivating feature of agent's cognitive character which contributes to her cognitive flourishing, and entails reliability in achieving truth (Zagzebski 1996). On the weakest notion, intellectual virtue is the proper function of agent's cognitive faculties, associated to the norm of reliability (Sosa 1991). There are views inbetween such as James Montmarquet's view (Montmarquet 1993) that virtue is a trait of cognitive character, but does not entail reliability, or Greco's view that it is a reliable faculty with minimal subjective perspective (Greco 2005).

Virtue epistemologists standardly divide into two camps, depending on how they understand the nature of intellectual virtues. The division depends on how close virtue epistemologists take intellectual virtues to be to moral virtues. One line is defended by Ernest Sosa and John Greco, who identify intellectual virtues with normatively characterized cognitive faculties: reason, intuition, memory, perception, and induction. I will call this view '*the faculty view*'. This line takes intellectual virtues to be independent, and different from moral virtues. The second line is proposed, and upheld by Linda Zagzebski, Lorraine Code, Christopher Hookway, James Montmarquet, etc. According to this line, intellectual virtues are traits (or qualities) of our cognitive capacities (like impartiality, intellectual honesty, trustworthiness etc). They are conscious and achieved personal disposition that motivate one in reaching the truth. This line treats intellectual virtues in the manner completely analogous

to moral virtues. Linda Zagzebski, who is the most radical proponent of the view, for instance, argues that intellectual virtues are best understood as a subset of moral virtues. Intellectual virtues like impartiality, intellectual honesty, intellectual courage, open-mindedness, trustworthiness etc. are on a par with moral virtues like courage, piety, magnanimity, generosity etc. I will call this second line '*the morality-centred view'*.

The opposition between the two lines draws from two relevant Ancient Greek uses of the word *arête* (virtue): virtue as a trait of character vs. virtue as a standard of functioning (the word for function is *ergon* which literally means 'task' or 'work'). Certainly the first use of the term *arête*: virtue as a trait of character, associated to Aristotle's theory of moral virtues, is much closer to the common (contemporary) intuitions of what a virtue is. This is so, because intuitively, virtue is primarily linked to the right *motivation* for good or truth. Julia Annas (2003) provides a good explanation of this intuition. The virtuous man has two goals (clearly explicated by the Stoics): the ultimate goal of virtue is virtuous life as a whole (happy life); and the concrete goal is the goal of particular virtuous action. Reaching the ultimate goal is a question of right motivation and is up to the person, whereas reaching the concrete goal is not entirely up to the person, but can be due to moral and epistemic luck. To illustrate this statement, Annas gives the example of Socrates, whose defence in the court fails in achieving its (concrete) aim of producing acquittal, but we still consider it virtuous (in terms of achieving his ultimate goal). If this explanation is correct, virtue is an expression of what is up to the person, and only reaching the ultimate goal seems crucial for virtue. Therefore, having a virtue is a question of right motivation and its fulfilment in a long-term period, which reveals a trait of character of a person. It looks unusual at a first sight to take cognitive faculties to be virtues. One reason for that is because having a reliable cognitive faculty is not up to the person. So, it does not look as if we should praise a person for having reliable reasoning, or memory, etc

There is, however, another relevant Ancient Greek notion of virtue, which takes virtue to be a broader concept than just a trait of character. This broader notion marks the characteristic excellence of a given class of things (drawn from its standard of functioning) over other classes of things, such as the excellence of knives in cutting things, or the excellence of eyes in seeing. A passage from Plato's *Republic* is considered representative of this particular use of the term, but we also encounter the same use in Aristotle, as I mentioned. The relevant passage is from Book I, 352e-353d, where Socrates talks to Glaucon and Thrasymachus:

> Is it possible to see with anything other than eyes?
> Certainly not.
> Or to hear with anything other than ears?
> No.
> Then, we are right to say that seeing and hearing are the functions of eyes and ears?
> Of course.
> What about this? Could you use a dagger or a carving knife or lots of other things in pruning a vine?
> Of course.
> But wouldn't you do a finer job with a pruning knife designed for the purpose than with anything else?
> You would.
> Then shall we take pruning to be its function?
> Yes.
> Now, I think you'll understand what I was asking earlier when I asked whether the function of each thing is what it alone can do or what it does better than anything else.
> I understand, and I think that this is the function of each.
> All right. Does each thing to which a particular function is assigned also have a virtue? Let's go over the same ground again. We say that eyes have some function?
> They do.
> So there is also a virtue of eyes?
> There is.
> And ears have a function?
> Yes.
> So there is also a virtue of ears?
> There is.
> And all other things are the same, aren't they?
> They are.
> And could eyes perform their function if they lacked their peculiar virtue and had the vice instead?
> How could they, for don't you mean if they had blindness instead of sight?

> Whatever their virtue is, for I'm now asking about that but about whether anything that has a function performs it well by means of its own peculiar virtue and badly by means of its vice?
> That's true, it does.
> So ears, too, deprived of their own virtue, perform their function badly?
> That's right.
> And the same could be said about anything else?
> So it seems.

Such passages affirm the theory that there is an interconnection between good, virtue, and function (*ergon*). They claim that the value terms *arête* (excellence) and its opposite *kakia* (evil) are ascribable to each thing which has a function. If a thing implements its function well, we ascribe to it excellence, if it implements it badly, we ascribe it a defect or imperfection.

Sosa and Greco apply the functional model of virtues to our cognitive faculties. Their reliability is the main source of excellence. Reliability depends upon two components: normal functioning of our cognitive faculties (where normality is cashed out in terms of the standard cognitive system of the species) and collaboration from appropriate circumstances. Our faculties have their natural functions to produce truth and to avoid error in a given thematic field and in given circumstances. As a knife will cut a pineapple if it is not a plastic toy and if there is no wire in the pineapple, *S*'s perceptual faculty will acquaint *S* with the fact that there is a white cup in front of *S* if *S* is not overwhelmingly short-sighted and if there is no special trick, such as a hologram of a cup instead of a real cup in front of the perceiver. Accordingly, a faculty would be vicious if a specific perceptual defect affected the capacity.

To generalize, a common point between the two notions of virtue concerns an agreement that virtue unifies internal and external components: motivation toward certain goal, and success in achieving that goal. The difference consists in the stress put on each of these components. Morality-centered approaches take the motivational component which

entails reflective internal control to be determinate of intellectual virtues. The proponents of this line claim that intellectual virtues can be understood in terms of exercising habituated and sustained motivation for truth in terms of doxastic responsibility. The faculty view, on the other hand, stresses the presence of a reliable method for achieving truth, and assigns a more modest role to the motivational component. For instance, it is not clear whether a motivational component has any role to play in Sosa's purely functional theory of virtues. If there is any sort of agential control on his theory, it amounts to implicitly keeping track of the reliability of one's sources of knowledge. According to John Greco, the motivational component of one's intellectual virtues is expressed by an agent's trying to sustain a well integrated cognitive character. The agent is properly motivated, for instance, if she strives to integrate her occasional strange and fleeting belief forming dispositions with the rest of her already integrated cognitive dispositions. If such integration is impossible, the agent is bound to avoid forming beliefs on the basis of strange and fleeting processes.

1.5 The role of intellectual virtues in an account of knowledge

A natural question to ask at this point is which of the two notions of virtue can play a better role in an account of knowledge? A *prima facie* objection against the morality-based view is that beliefs accepted out of right motivation (or responsibly accepted ones) do not necessarily yield truth, and truth is indispensable for knowledge. As I said above, the morality-based view of virtues is shaped on the model of happiness: what is important is the acquisition of ultimate goal, which can dispense with the acquisition of immediate goal. However, every particular instance of knowledge entail truth. So the question is how then the *motivation* for truth is helpful in explaining knowledge?

Linda Zagzebski (1999, pp. 92- 116) tries to overcome this problem by defining knowledge as a result of an *act* of intellectual virtue. Her definition of 'act of intellectual

virtue' is complex. Only when an epistemic act *actually achieves truth* and *the subject of this act* has *a stable disposition* to exercise acts of this sort, the act is an *act of intellectual virtue.* An act by which one achieves truth but lacks stable disposition is not an act out of intellectual virtue, and an act by which one does not succeed to achieve truth, but has stable disposition to achieve truth in the range of circumstances, is not an act out of intellectual virtue. According to Zagzebski's definition:

> [A]n act of intellectual virtue *A* is an act that arises from the motivational component of *A*, is something a person with virtue *A* would (probably) do in the circumstances, is successful in achieving the end of the *A*'s motivation, and is such that the agent acquires a true belief (cognitive contact with reality) through these features of the act. (1996, 271–1)

This definition of 'act of intellectual virtue' is not a kind of definition that we would expect from a proponent of the morality-based view, because it violates the main intuition behind the view. It requires from a virtue to achieve its immediate goal, and not only to achieve its long-term goal. But as Annas rightly points out, it is simply in the nature of moral virtue not to always achieve their immediate goal.

Annas (2003) indeed criticizes Zagzebski's definition of knowledge. On the one hand, she argues, if intellectual virtue is a subspecies of moral virtue in Zagzebski's style, it must have a Stoic structure, according to which the ultimate goal, namely right motivation is basic in the definition of knowledge. On the other hand, Zagzebski considers concrete success as a key characteristic of the act of intellectual virtue, and thereby shows that in achieving truth we need precisely what *virtue cannot supply.* Annas's line of thought suggests that Zagzebski's attempt to solve the epistemic problem for the morality-based view is unsuccessful. If Annas is right, the morality-based theorists have a serious problem of explaining how virtues are indispensable for knowledge, because it is simply not in the nature of their understanding of intellectual virtue that every act of virtue entails truth.

The faculty view, on the other hand, seems to imply the opposite problem, namely it cannot straightforwardly account for the motivational component of our intellectual virtues. Sosa has been criticized, for instance, for not having exploited sufficiently the potential that virtue theories offer (Montmarquet 1993). However, given the functionalist notion of virtue that they embrace, this is not supposed to be a problem for the view. There is a further question, though, whether such a weak notion of virtue is sufficient to overcome the weaknesses of generic reliabilism.

Intellectual virtues have been involved in the epistemological enterprise for two main reasons. One such respect is supplying generic reliabilism with a normative component and thus avoiding its tendency towards naturalized epistemology. Virtues seen as proper functions of our cognitive faculties are sufficient to do that. They set a standard, and therefore a norm for the source of knowledge. Also, a central problem that virtue theories are meant to address is the so called "value problem of knowledge". The value problem refers to, and inherits Plato's concern of why knowledge is more valuable than mere true belief. This particular form of the problem is also known as "the primary value problem". A contemporary version of the problem which has sprung in the post-Gettier era is the so called "secondary value problem" which relates to the concern why knowledge is more valuable than a proper subset of its parts, for instance, why is it more valuable than true and justified belief. The secondary value problem grew into the so called "tertiary value problem": why is knowledge more valuable than the proper subset of its parts not just in degree, but also in kind (Pritchard 2007: §2)?

The answer that virtue epistemologists give to the primary value problem is that knowledge is more valuable than true belief because it results from praiseworthy characteristics of cognitive agent. They answer the secondary and the tertiary value problems by claiming that knowledge is more valuable than justified true belief both in degree and in

kind, because it is not just belief that is true and virtuously formed, but it is belief that is true *because* virtuously formed. It seems that the faculty approach is in a position to address the whole range of value problems. Knowledge is more valuable than true belief because it results from reliable faculties as opposed to mere guessing. It is more valuable than justified true belief both in degree and in kind, because it is not just a belief that is true and reliably formed, but it is belief that is true *because* reliably formed. So, in order for a theory to solve the value problem, it need not understand credit as something more that attributability to reliable faculties. Of course, this is only a brief defense which requires further elaboration. Chapter 4 of this thesis is partly dedicated to this problematic.

What about the persistent problem of the Clairvoyant Norman? Can we address this problem just on the basis of a functionalist understanding of virtue? Intellectual virtue is meant to supply the reliabilist criterion of knowledge with a minimal internalist or subjective perspective over and above the requirement of *de facto* reliability. Greco (2003), and Daniel Breyer and John Greco (2008) have provided a story of how subjective certainty comes into the purely reliabilist picture. Greco (2002, 2003) distinguishes between subjective and objective justification. Objective justification is measured in terms of *de facto* reliability of one's faculty, and subjective justification is understood in terms of 'cognitive integration'. Cognitive integration is the integration of our faculties on the first level which entails minimal control and responsibility from the agent, i.e. control from one's virtuous epistemic character. The idea is that when the agent holds a default status of being motivated to believe truth, a faculty *f* is getting cognitively integrated with the rest of the agent's faculties in virtue of that motivation. This integration is fulfilled because the agent is motivated to believe the truth, although the integration itself can be sub-personally implemented. As Breyer and Greco (2008, p. 174) put it: "Here "motivated to believe the truth" is meant to pick out a default status that believers are typically in when forming their beliefs." The cognizer is subjectively

justified only when she forms beliefs out of her well-integrated faculties. So, Norman does not know until he satisfies the subjective requirement of cognitive integration.[11]

Greco (in Greco and Turri 2010) claims that he offers a mixed theory of virtue – a combination of functionalism, and morality based approach, where the moral component is minimal. However, we can ask why cognitive integration cannot be treated as a proper function? Breyer and Greco state that the result of integration brings to mind a feeling of "belief ownership". In other words, when a belief is formed through a cognitively integrated faculty it is issued as one's own, and not as alien. When one has this feeling, one is entitled to believe. Thus, the theory of cognitive integration arguably supplies the functional theory of virtue with a minimal subjective component. But why is this a moral element? Even if our faculties have to be exercised with a feeling of appropriate certainty, this feeling is based on, results from, and is to be credited to a properly functioning cognitive system, and not to an agent's conscious effort. In chapter 3, 4, and 5, I clam that the notion of credit is integrated into the notion of intelligent ability, an ability exercised with the right propensities invested by a cognitive system. I provide some hints, but I do not develop a full-fledged answer to the Clairvoyant case along these lines.

1.6 Are virtues really necessary for knowledge?

As we saw above, one important line in virtue epistemology has been following an ancient Socratic and Stoic tradition according to which virtues are like skills, or capacities. In what follows, I want to propose an even more ability-centred interpretation of the main epistemic norm. I will try to show that all benefits that can be drawn from the notion of intellectual virtue for an epistemological theory can be drawn from the notion of intelligent

[11] Bernecker (2008) expresses doubt that Greco's theory of cognitive integration manages to accommodate the internalist element required by BonJour for a tenable solution of the Clairvoyant case. Breyer and Greco reply that 'subjective' does not mean 'internalist', and that a proper subjective element is sufficient to solve the problem.

ability. Abilities are normally identified teleologically (see, for example Millikan 2000, pp. 59-61) with a view to their task, and to what would count as success in deploying them. Sometimes, they are seen as merely instrumental, and their goal is described independently. For instance, my aim in the Chinese restaurant might be to satisfy my hunger, and the way I eat my meal is merely instrumental. Alternatively, my aim might be more sophisticated, namely to satisfy my hunger in the way prescribed by Chinese eating culture, and in this case the skilful use of sticks is not merely instrumental, but is part of the goal. Our ordinary cognition suggests the luxury of the later case is redundant. I shall argue that at least some part of our knowledge results from our attempts to get to the truth in the most immediate way, and not from our attempts to get to the truth in the most sophisticated (virtuous) way. Without explicitly thinking of the latter, we achieve the former.

1.6.1 Arguments against the necessity of intellectual virtues for knowledge

I take it that a basic commitment shared by all virtue epistemologists is that intellectual agents are the primary source of epistemic value and the primary focus of epistemic evaluation. If we are to embrace a virtue epistemology, we also have to embrace the idea that the primary epistemic value is ascribable to an agent's epistemic character, at a minimum to one's proper functions. Virtue as a function is still an excellence; it is a *stable* disposition to achieve truth in certain domain. For instance, if I am a virtuous thinker in a particular domain, I must be a reliable thinker in that domain. If knowledge essentially depends on virtue (is caused by virtue) then, knowledge essentially depends on (is caused by) such an excellence, or in a stronger sense by one's virtuous epistemic character. Hereafter, I will be interested in the question is virtue really necessary for knowledge, and I will claim that it is not.

Here is my first argument against the necessity of virtues for knowledge. Knowledge is modular – some people are better in visual recognition than in mathematical calculations. If I am generally bad in mathematical calculations, I do not plausibly exercise a virtue when it comes to the domain of math. My mathematical reasoning is generally deficient – I am prone to making mistakes. It would be highly inappropriate to say that I exercise a virtue of reasoning when it comes to mathematics. Nevertheless, I can have mathematical knowledge. I can do simple calculations and even more complex ones if I try hard. From this example, it follows that virtuous character is not necessary for knowledge. Therefore virtue is not needed for knowledge.

A possible reply to my argument is to claim that virtue is cognitive integration (after Breyer and Greco 2008). The exercise of virtue entails that whatever belief is formed, it must be cognitively integrated with the rest of one's knowledge. On this theory one does not need to be good in all cognitive domains, but just to be a good cognitive integrator. From this viewpoint, a belief results from a virtuous character and therefore amounts to knowledge whenever it is true and it is cognitively integrated with the rest of one's beliefs at least in the relevant domain. Seen in this way, virtue can be perhaps reestablished as a necessary condition for knowledge.

My answer to this reply is that cognitive integration is often sub-personal. But if it is sub-personal, it is not creditable to the agent. An agent is unaware of many things going on the sub-personal level, e.g. in her brain. Therefore, cognitive integration is not a virtue. Yet, if you don't find this answer compelling, think about the following. Only some beliefs or instances of knowledge are relevant for cognitive integration when I do math. For instance, my perceptual beliefs – such as "that tree has yellow leaves" – are not relevant for cognitive integration when I do math. The relevant beliefs for cognitive integration include other mathematical beliefs – such as the rules of subtraction, division, and other laws of

mathematics. But this is what I am bad at – to take into account in a reliable way the right set of relevant beliefs when I have to solve a mathematical problem. Therefore, I am not a good integrator in the domain of math. But I can still have mathematical knowledge. I am simply a bad integrator, but sometimes, I do take the relevant set of rules into account, and I achieve knowledge. Therefore, knowledge does not essentially depend on the exercise of virtues. Yet it is still essentially depends on the exercise of some intelligent skill of mine, which in the concrete example is far from being brilliantly intelligent.

Another objection against the necessity of virtue for knowledge taken as cognitive integration can be generated in the following way. There is a conceptual possibility that I know just one proposition *A*: "This is A". Suppose I have no other beliefs with which I can integrate *A*. If cognitive integration is necessary for knowledge, it is unexplainable how I can know *A*.

Next, it is clear that one's agency and conscious motivation need not always be involved in acquiring abilities and knowledge. One can even acquire an ability contrary to one's will. Take, for instance, a possessive husband who does not want to learn how to cook, so in order for his wife to stay at home and cook for him. However, being around his wife at home, and observing how she cooks, he eventually learns how to cook. Sometimes he has to prepare the products, to stir the meal, and practically to do the whole dish while his wife is occupied with other affairs. He learns how to cook without any conscious motivation.

A third argument, commonly raised against the necessity of virtue for knowledge, is an argument from perception. Perception is involuntary, and automatic. Hence even if perception is a subject of refinement, perceptual refinements are subconscious. Then, perception is not a virtue creditable to the agent. The agent often has no idea how she is able to perceive.

Perception is rather creditable to the agent's cognitive system. I examine the argument in more detail in chapter 5.

Finally, although intellectual virtues (taken as reliable faculties or stable dispositions to reach success in field F in circumstances C) seem not to be necessary for having animal knowledge, they might still be a necessary condition for having reflective knowledge in Sosa's sense. If I am to reflectively know that my mathematical faculty is reliable, it has to be reliable. Reflective knowledge is a reflection on the reliability of our faculties. If I am to have reflective knowledge about my condition and the reliability of my faculties, at least some of them have to be virtuous (they have to be stable reliable dispositions). So, having virtuous faculties is a necessary condition for reflective knowledge. However the account of knowledge provided here is limited to the basic first-level knowledge.

These considerations yield the conclusion that we cannot take for granted that agents and their characters are the primary focus of epistemic evaluation. Achieving knowledge is a complicated process of which the agent can be mostly unaware. So, it is not the agent who has to be credited for her reliable faculties. Therefore, reliable faculties are not virtues creditable to the agent. On the other hand, it does not seem that the agent's epistemic character is essentially in charge of knowledge. Virtues are not minimally necessary for knowledge, although they might contribute, and enhance our ways to knowledge in general, and they might be necessary for some particular instances of knowledge.

However, there is much to be learnt from virtue epistemology, and in particular from Sosa's and Greco's virtue reliabilism. In the following chapters, I defend a weaker version of the theory by trying to suggest new aspects, and possible benefits of it. For the reasons just outlined, I avoid the vocabulary of virtues, and adopt instead the term 'intelligent abilities'.

CHAPTER 2: THE PRIMENESS OF APT BELIEF

In the previous chapter, I tried to explain the motivation for virtue epistemology, and the significant progress that it makes in solving the major problems related to the preceding views. My presentation was focused on the faculty view. In this chapter, I am going to explore the potential of providing a metaphysical analysis of knowledge based on the faculty view. Ernest Sosa (2007, pp. 22-44) offers such a definition in the form: *S* knows that *p iff S* believes aptly that *p*, where apt means accurate because adroit, or true because competent.

It is a common-sense wisdom known by any sports coach or dance teacher that exercise of (interesting) abilities involves protracted and often subtle coordination of "inner" and "outer" elements. A good instance is tango. The tango leader has to coordinate his body movement not only with the rhythm in accordance with various leading techniques, and his own creative ideas, but at the same time he has to follow the circle of other dancing couples, and to pay attention for the whole couple to move safe. All this requires a complicated coordination at the same time. Timothy Williamson has offered an ingenious technical characterization of such coordination in terms of what he calls "primeness".

In this chapter, I claim that knowledge is like that – it results from the exercise of an intelligent ability trained to achieve a sufficiently flexible coordination with its object, and it requires such coordination. My aim here is to extend Williamson's defense of primeness of knowledge to practical abilities, and to show that the notion of knowledge as success from ability is a prime notion.

The chapter has three main parts. In the first part, I discuss what I call "the conjunctive view of knowledge"- the view that knowledge is a merely conjoined satisfaction of certain

internal and certain external conditions. The classical definition of knowledge as true and justified belief is an example of the conjunctive view. Then, I discuss the Gettier problem, as well as some attempts to resolve it. In the next part, I sketch two alternatives of the conjunctive proposal: "radical anti-conjunctivism" – the view that rejects the possibility of analysis of knowledge, and "mild anti-conjunctivism", according to which knowledge can be analyzed in terms of a relational component between the internal and the external. In the third part, I argue in favor of the second position. In particular, I defend the view that knowledge is success from ability, or in short (KSA) by showing that it is resistant to Gettier counterexamples and to the primeness argument. In this context, I present and discuss Sosa's view of knowledge as apt belief, which is foundational of KSA.

2.1 Conjunctive view of knowledge

Epistemology has been traditionally engaged with explaining knowledge, although nowadays, and in particular under the influence of virtue epistemology, some epistemologists think that knowledge is not the central epistemological notion[12]. I am a traditional epistemologist in this sense: I think that knowledge is central to our epistemic activity. So, my project is going to deal with the notion of knowledge.

Together with substantial number of sophisticated discoveries related to the nature or meaning of the concept of knowledge, contemporary epistemology, of course, offers a number of debates and disagreements as we saw in the previous chapter. However, the crucial role of two features of knowledge has always been for the most part beyond dispute. First, that knowledge entails truth and second, that if an agent is to know, she has to contribute to grasping the truth as opposed to merely guessing it right. Technically speaking – she must grasp the truth in a non-accidental manner.

[12] Most prominently: Kvanvig (1992) (2003), Roberts and Wood (2007), and also to some extent Riggs (2003).

Traditional approaches have been preoccupied with searching for necessary and sufficient conditions for knowledge that involve these two elements. Truth and the non-accidentality component (generally "justification") have been seen as two independently determined components, belonging accordingly to the external and the internal realms that, taken together, give us necessary and sufficient conditions for knowledge. Traditional approaches tell us that these two components must be jointly satisfied in order to give us knowledge. I shall call this "the conjunctive approach" to stress the merely external nature of the way in which the two components are put together. The stress is on the assumption that each can be specified independently of the other and merely aggregated together in order to characterize knowledge. An example of that is the classical definition of knowledge as true *and* justified belief, following Plato's passage from "Theaetetus" (201d-210a); and another example is arguably its latter branch, the reliabilist account of knowledge.

2.2 The Gettier problem

Such merely conjunctive approaches have been famously undermined about fifty years ago. Edmund Gettier has taught us the lesson that if one has true *and* justified belief, one is not automatically provided with knowledge. Therefore, the crucial components, truth and the agent's contribution, cannot conjointly give us an explanation of knowledge. The more recent fake barn example, in Goldman's version, became similarly famous along these lines. Finally, Timothy Williamson made an even more precise contribution along the same lines: he attacks the conjunctive analysis of knowledge by claiming that knowledge is a prime mental state which is not decomposable into internal and external components. We will discuss all the three objections.

In his paper "Is Justified True Belief Knowledge?" Gettier (1963) presents two effective counterexamples to the analysis of knowledge as true and justified belief.[13] Here is my own Gettier-type example. Suppose that you and I play chess. At some point of the game, I realize that your next step will be knight on C4. I see clear reasons for that, namely that the move will enable you to take my bishop. So, I believe:

B: Next, you will play knight on C4.

Your next step is actually knight on C4. However, you have not even thought of the possibility to take my bishop. In fact, what you are trying to do is to take my queen in five steps. Did I then know that you will play knight on C4? It seems that I did not, despite the fact that my belief is true, and justified. My belief is justified, but it is justified for the wrong reason. However, it hits the truth by sheer accident. Cases like that constitute a problem, because there is undesirable luck involved in them.

There are two general strategies in the latter literature to avoid Gettier counterexamples. One of them appeals to the need of a fourth independent condition in addition to the original three. The idea is to *add* a further condition for knowledge that will block the Gettier cases. The resulting account is known as the TJB+G account. The alternative approach aims at finding a link between truth and justification which cannot be broken, in order to block the Gettier counterexamples.

Initially, it was thought that the Gettier cases arise when an accidentally true belief is inferred from a false premise. In our case, this is the premise that you want to take my bishop. One theory which takes this characterization into account was offered by Lehrer and Paxson

---

[13]One of the original Gettier examples can be summarized as follows. Smith sees his acquaintance John driving a new Ford. He knows that John has always owned Fords beforehand, so he forms the following belief on the basis of good evidence: 1. *Jones owns a Ford*. That belief, however, turns out to be false. Now, suppose that Smith infers from 1 the following further proposition: 2. *Either John owns a Ford or Brown is in Barcelona*. Smith recognizes the legitimate entailment, so he is justified in believing that proposition too. Also, as a matter of fact, Brown turns out to be in Barcelona. So, 2 is true and Smith is justified in believing it. However, it is commonly agreed that Smith does not know 2.

(1969). It appeals to a requirement for absence of a defeater. Knowledge is taken to be *undefeated justified true belief* — which is to say that a justified true belief counts as knowledge if and only if it is also the case that there is no further fact that, had the subject known it, it would have defeated her present belief. Some critics of the view such as Marshall Swain (1974) have argued that the notion of a *defeater* cannot be made precise enough to rule out the Gettier cases without also ruling out legitimate cases of knowledge. Another problem related to the non-defeater theory concerns our beliefs based on induction: in order to preserve the large portion of knowledge based on induction, the proponent of the theory needs to specify the relevant notion of defeater since beliefs based on induction are always open to possible defeaters no matter how far-fetched they are. A third problem is that it looks ad hoc, because it is constructed to solve the Gettier problem, and it does not provide a positive story about the essence of knowledge. After all, the insight behind the Lehrer-Paxson theory does not add much to the very discovery made by Gettier, namely that true and justified belief can be defeated.

In fact, the characterization of the Gettier case as inference from false premise is by itself problematic. The problem is that it is also possible to generate Gettier cases about beliefs that are directly justified on the basis of experience, and not inferred from other beliefs. A relevant example is provided by Matthias Steup (2008) in which a person seems to see a dog eight yards away from her, and so forms a belief that there is a dog there. In fact, what she sees is a Japanese toy, perfectly indistinguishable from a real dog. The toy is put there by the manufacturer to test the public reaction. Somewhere behind the dog, though, there is a real dog. So, the person happens to have a true belief which is arguably directly justified on the basis of visual perception, but she lacks knowledge. [14]

---

[14]This is a problem for Lehrer and Paxson, because they believe that only non-basic beliefs can be defeated.

Let us look at an alternative explanation, according to which in the Gettier cases, the correlation between truth and justification is broken. The root of this strategy is to be found in Goldman (1967) who suggested the addition of a *causal* condition: a subject's belief is justified, only if the truth of a belief has *caused* the subject to have that belief (in an appropriate way); and for a justified true belief to count as knowledge, the subject must also be able to "correctly reconstruct" (mentally) that causal chain. This account is supposed to block examples of Gettier beliefs justified directly on the basis of experience, because these are cases where the causal link from the object to the belief is broken. However, Goldman faces the difficulty of giving a principled explanation of how an "appropriate" causal relationship differs from an "inappropriate" one without the circularity of saying that the appropriate sort of causal relationship is the knowledge-producing one.

Later on, Goldman (1976) and Dretske (1971), among others, proposed the substitution of: 'caused by the fact p' with 'produced in a reliable way'. One big difference is that on the first notion, there is no false belief that satisfies the condition, and on the second, there is. Ultimately, the approach appeared to be ineffective against the Gettier cases, because Gettier beliefs are indeed usually formed by reliable methods.

It took some time for philosophers to realize that any analysis of knowledge that allows for truth and justification to go apart in a concrete case is vulnerable to Gettier-style counterexamples. In her article “What is Knowledge”, Zagzebski (1999, p. 101) claims that when two different chains lead to justification and truth, the account is always vulnerable to Gettier type cases. Gettier cases are cases of double luck, she notices: bad luck is canceled by good luck. In our example, it is a piece of bad luck that I judge about your next step on the chessboard by using false evidence, and then a piece of good luck that you play accordingly. It is not surprising that the reliabilist account fails in the face of the Gettier examples, since a reliable method does not determine truth in every particular instance of belief formation. To

use again our chess play example, relying on the kind of evidence that I relied upon in predicting your next move is a reliable method of belief formation, but this time *it* did not lead me to truth. Gettier cases can be generated when the second condition makes truth only probable, but does not entail it. Linda Zagzebski thinks that the reliability account says too little, because it does not say what the proper conceptual connection between the two components is. Zagzebski rightly notices that normative properties of beliefs cannot provide the necessary link that we look for, so she suggests that we look at normative properties of persons instead.

The moral of the Gettier debate is that if a definition of knowledge is to be proposed, it should not be a mere summation of the two elements: truth and the justifier. It is clear that any attempt to build up an account of knowledge by conjoining a set of independent conditions was misguided from the outset.

2.3 Non-conjunctive proposals

So far, we saw that the conjunctive accounts of knowledge are nor satisfactory. Now what are the alternatives? Despite the objections against Goldman's causal theory, and the reliabilist account of knowledge, the strategy of looking for a link between the two components could still be on the right track. The only problem, as we saw it, is that the relational component we are searching for has to grant a perfect correlation, which is infallible in particular instances. But is it possible to find such a relation without circularity? The answer to this question determines the possibility of preserving an analysis of knowledge. One thing is certain - the only direction to go is to deny traditional conjunctivism. And then the question is which way to take from here? There are two camps along the "anti-conjunctivist" direction. I will call them "mild" and "strong" or "radical" anti-conjunctivism.

2.3.1 Radical anti-conjunctivism

Those who have adopted strong anti-conjunctivism embrace skepticism about the possibility of analysis of knowledge. They generally argue that epistemological terms like justification, evidence, certainty, etc. should be analyzed in terms of a primitive notion of *knowledge*, rather than vice versa. Knowledge is understood as not decomposable into internal and external components, such as justification and truth, and hence as being unanalyzable. So, I will call "radical anti-conjunctivism" the thesis that knowledge is unanalyzable.

Timothy Williamson's view of knowledge as a mental state is such a radical anti-conjunctivist view. Williamson (2000, p. 49) denies that knowledge consists of a satisfaction of independent internal and external conditions. He takes knowledge to be a mental state which extends to, and encompasses the external object. Knowledge, Williamson thinks, is the actual union of an agent's cognitive state and the fact: it is a mental state not just entailing a fact, but in a stronger sense – reaching out and encompassing a fact. This applies naturally to all derivatives of knowledge such as perception, remembering, recognizing etc. All these are mental states that are not decomposable into internal and external components. The central argument that Williamson offers in favour of this view is the so called "primeness argument". I explain the argument in the next paragraph.

2.3.1.1 The primeness argument

Tim Williamson (2000, pp. 65-93) provides one major argument in support of the thesis that knowledge is a mental state: "the primeness argument". Like the Gettier cases, the primeness argument shows that a conjunction of internal and external conditions is insufficient for knowledge. The primeness argument is a *reducio ad absurdum* of the thesis that knowledge can be decomposed into internal (narrow) and external (environmental) components. The internal component is the component of believing that p, or justifiably

believing that *p* in a narrow sense. The external component is the presence of *p*. Now, suppose that knowledge is just a conjunction of such narrow and environmental conditions. Williamson invites us to consider three different cases: α, β, and γ. Suppose that there is water in front of a subject *S*. In α, *S* sees water with her right eye only and there is a device which simulates water in front of her left eye, but she has a head injury preventing her from processing the input from that eye. The state β is the reverse of α: *S* sees water with her left eye only, because there is a water-simulating devise in front of her right eye, and she has a head injury preventing her from processing the input from that eye. Finally, the state γ is internally like α, and externally like β. It is a situation where the left eye in front of which there is real water does not process anything, its brain route is injured, and the eye in front of which there is a water-simulating device produces the corresponding "narrow" content of water. The moral of the story is that in γ, *S* does not see, although the purported internal and external conditions for seeing are satisfied each one separately and conjointly. Hence, seeing is not a conjunction of internal and external components. This generalizes to knowledge: knowledge, as well as seeing is a prime state.

Tim Willamson (2000, pp. 2-5, and pp. 27-33) makes a separate point by claiming that it's not only that the state of 'knowing' is prime, but also the concept of knowledge is unanalyzable into sufficient and necessary conditions. He thinks that there might be necessary conditions for knowledge such as belief or justification, there might also be sufficient conditions for knowledge such as seeing, but no conjunction of necessary and sufficient conditions can yield a satisfactory analysis of knowledge. Indeed, he argues, the concept of knowledge has conceptual priority to concepts like belief and justification, rather than vice versa. However, the claim of unanalyzability of knowledge does not follow straight from the primeness argument, so Williamson argues separately for it.

One might think that what is missing in the primeness argument is an adequate connection between the internal and the external components. The theories, according to which knowledge is decomposable into internal and external bits, usually take it that the state of knowledge is internal, and that there are some external states of affairs that have to be present, together *with some link between the two components*. Perhaps an adequate condition for the correlation between the two components would provide the sufficient condition for knowledge.

Take the classical internalist theory of justification, according to which knowledge is justified true belief. The epistemological link between truth and the internal state of knowledge is displayed by some reason that the subject possesses for believing that the object in front of her. The primeness argument works against this account in the same way as the Gettier examples do. In case γ, *S* possesses a reason to belief that there is water in front of her, and there is actually water in front of her. But she has no knowledge. So, justification of that sort is too weak to provide a sufficient condition for knowledge.

Next, take an externalist theory of justification according to which the subject has to be reliable, and there should be a causality type relation from the object to her belief. In this case, we can always modify Williamson's example to make the state causally over-determined by a bad cause in a way that would make the causal chain from the object to the subject deviant. For instance, we can add a demon, who usually switches the button of the water simulating device on when there is water in front of the subject, and treat γ as a case in he does so. Then, there would be a causal chain going from the external bit to the internal bit: the presence of water causes the demon to switch the water simulating device which causes the cognizer's belief that there is water in front of her. So, the condition is satisfied. But the cognizer does not seem to have knowledge in that case either.

So, Williamson's first step in the argument from unanalyzability is a pessimistic induction from the existing attempts to analyze knowledge into separate internal and external components and a link between them, to the conclusion that no such analysis is possible without being vulnerable to counterexamples, or circular.

There are two ways for the proponent of the analysis to go from here. The first way is to argue that some of the existing definitions are sufficiently good approximations to indicate strongly that a further refinement on similar lines will eventually succeed. But, Williamson continues, 'a sufficiently good approximation' does not entail conceptual priority of that approximation to the original concept.

The second way is to try to find a link between the internal and the external components so strong that it would block counterexamples. Williamson suggests that such a link cannot be spelled out without relying on the notion of knowledge. If any of the above theories has to fix the relational component in order to avoid counterexamples, the relational component will amount to "in the way which yields knowledge", and therefore would be circular. In order for it to be non-circular, Williamson continues, we should be able to think of the components of the definition without even implicitly thinking of knowledge.

Now, let me summarize the whole argument. First, the primeness argument shows that mere conjunction of internal and external components is not sufficient for knowledge. Then, Williamson argues that an adequate connection between the conjuncts is parasitic on knowledge. Therefore we can conclude that reductive analysis of knowledge is hardly possible.

Despite the fact that the situation at this point does not look rosy for the proponent of the analysis, let us consider one last option. This is the option, according to which knowledge is accepted to be a broad mental state, and at the same time it is claimed to be analyzable. Remember the line that virtue epistemologists follow in solving the tertiary value problem:

the problem why is knowledge more valuable than a proper subset of its parts. Virtue epistemologists claim that knowledge is more valuable than justified true belief in kind because it is not just belief that is true and virtuously formed, but it is belief that is true *because* virtuously formed. Now, my suggestion is that reaching truth because of intelligent ability can be conceived as a broad mental state. In cases of knowledge, the exercise of intelligent ability puts one in direct contact with the objects of knowledge. In a sense, our abilities enable us to reach out, and encompass the object. However, they do so in a fallible way, i.e. only in the presence of appropriate circumstances, not in a way which yields knowledge. The claim is then that when truth is reached via intelligent ability, the final result is knowledge.

The safety belt for such a definition of knowledge comes in two parts. First, it does not look to be circular, at least badly circular. The way virtue brings in truth is not a way which presupposes knowledge. Although the exercise of ability is a success notion, as I argue in chapter 5, it is such because it requires the presence of appropriate circumstances, not because of the way in which the truth is being grasped. For that way is fallible. Ability is just a disposition to reach truth, but it can guide one to incorrect beliefs too.

Second, since in 'true *because* competently formed' 'because' is causal, the first intuitive counterexample is a threat of deviant causal chains, or overdetermination. In these cases, the ability would be in the beginning of such a chain, the truth would be at the end of the chain, and a demon would be mixing inbetween. To use the example above, whenever a cognizer is to encounter water, a demon anticipates her potential belief formation, and switches a water-simulating device in front of her eyes at the time in which water is already present. In this case, it is the cognizer's ability that causes the demon to make her form a true belief whenever water is present. However, as I argue, an apt state is a broad state. Moreover, it entails the requirement that truth is captured by the cognizer. Hence, such a chain cannot be

generated, because it is incompatible with cognizer's subjective (apt) grasp of the truth. The demon helper deprives the cognizer of her own grasp of the situation, and therefore of the possibility to exercise her ability even if she has one. Since such a deviant causal chain cannot be generated, we have a reason to think that an analysis of knowledge of the suggested kind is defensible. In the next section I explore two such analyses.

2.4 Mild anti-conjunctivism

Let me call "mild anti-conjunctivism" the view that informative enough analysis capable of blocking the Gettier-like counter-examples is possible. I will present two attempts to analyze knowledge into internal and external components that are compatible with taking knowledge to be a broad mental state.

Linda Zagzebski (1999) proposes a version of such an analysis. She treats the problem of knowledge as an enterprise of finding a conceptual connection between truth and the sense in which knowledge is good, and she claims that this connection will provide us with the crucial element that we are looking for. Zagzebski thinks that the component which makes knowledge good is its origin in virtue. Virtue combines admirable internal state with external success. But what combines them in a really infallible way, she thinks, is the *act* of intellectual virtue.

Let us look at Zagzebski's particular idea in a little bit more detail. Zagzebski introduces the assumption that an act of intellectual virtue entails success and necessitates the implication from the truth- providing component to the truth of the belief. However, any such act requires right motivation (an internal component) and the presence of appropriate circumstances (an external component). Virtue has two components, according to Zagzebski: motivational component or the right emotion that guides one's action towards an end, and a success component – reliability in reaching that end. Motivational component entails only

having a disposition to a motive. It does not entail having the motive in every situation where the relevant circumstances are present. Reliability involved in the success component, on the other hand, does not entail success in every single instance. Hence, an act of virtue can be both motivated and reliably formed and yet deficient in some respect. It can fail to get the truth. This happens in the Gettier cases where the motive for belief formation is not responsible for the truth of the belief. The end must be reached *because* of the praiseworthy characteristics of the agent's virtue, the right motivation.

Let me remind you how Zagzebski defines 'act of intellectual virtue':

> An act is *an act of virtue A* if and only if it arises from the motivational component of A, is an act that persons with virtue A characteristically do in the circumstances, and is successful in bringing about the end of virtue A because of these features of the act." (1999, p. 108)

Hence, Zagzebski offers the following definition of propositional knowledge:

Knowledge is true belief out of an act of intellectual virtue.

My main objection to Zagzebski's analysis is that it is implausible to think that an act of virtue always entails success. I already addressed this objection in chapter 1. Here, I would like to explain why, for a similar reason, her analysis of knowledge looks stipulative. Zagzebski thinks that knowledge results from an act of intellectual virtue in the same way as a morally good action results from an exercise of moral virtue. Take the following practical syllogism:

1. Helping friends is good.

2. A friend is in need.

3. Help him!

And compare it with:

1. Knowing is good.

2. Acts of open-mindedness entail knowledge.

3. Act out of open-mindedness!

Despite the purported analogy between the domains of moral and cognition, one can notice the disanalogy between the second premises of the two syllogisms. Premise 2 of the second syllogism is crucial, but it cannot be taken for granted (like the premise 2 of the first syllogism) because of the following consideration. One can act open-mindedly in a completely well-tempered way by accepting one's opponent's claim, and yet happen to be wrong for some surprising reason. The reason why this is so, is because the possession of virtue can be measured through a sequence of unsuccessful acts. Imagine that I cheat you and make you believe that a situation is dangerous, in order to test if you are brave or not. If you act bravely, I presume, you would be motivated by your brave character, and you will pass my test. If I am not allowed to assume that your acts are acts of bravery in this case (since I cheated you, and the danger is unreal), how can I judge that you are brave?

A remaining dissatisfaction in Zagzebski's reply is that in the moral domain we are often prone to consider acts with catastrophic results as acts originating in virtue. A man who jumps in the water to safe a drowning child and gets eaten by a shark without reaching his goal is a convincing illustration. It is beyond doubt that he performed an act of moral virtue. Are intellectually virtuous acts different in this crucial respect from moral actions? It does not seem so, because an epistemic act can be justified without amounting to knowledge. And the justifier on that theory is nothing but virtuous motivation. To put the objection differently, it is unclear why a virtuous act would be an act of virtue only in proper circumstances, since it is the same trait of character that motivates acts that happen to be failures due to certain forms of luck that prevent success. Moreover, in order to preserve the parallel with moral virtues, Zagzebski is bound to claim that an unsuccessful attempt to save the drowning child is not an act of moral virtue.

Interestingly, Linda Zagzebski thinks that one does not need to be virtuous to be capable of virtuous acts. She wants to allow for those who are virtuous-in-training: e.g.

infants, or animals to be ascribed knowledge too. They would act in a way that a person with virtue would act. The problem with this suggestion is that the anti-accidentality condition for knowledge in the terms of the right motive becomes too vague. If the agent lacks virtue, but acts in a way in which the virtuous person would act, there is nothing to guarantee that she did not have the right motive by sheer accident. What makes her imitation of the virtuous person really virtuous? The simple answer seems to be that only if the person has the virtue, her act out of virtue would be non-accidentally virtuous.

There is a graver problem with Zagzebski's allover definition due to its implausibility. It looks from the above arguments that the sense in which Zagzebski uses "act of intellectual virtue" in the definition is a more limited technical sense including only those acts that are successful. Now, if it is legitimate to make this sort of assumption, we can save the day for any of the old definitions of knowledge too. For example, we can equally well distinguish a reliable act from a reliable method, and define a technical sense in which any reliable act entails truth. But this is going to be ad hoc. Given its general implausibility, Zagzebski's notion of act of intellectual virtue appears ad hoc too, and therefore it could not play the desired role in a definition of knowledge. In the next section, I explore Sosa and Greco's versions of mild anti-conjunctivism.

2.4.1 Knowledge as success from ability

What we learn from the above debate is that we should look for a more natural basis of describing truth reached through competence. In this section, I am going to present a principle which will be central to my debate until the end of the thesis. The principle is called "knowledge as success from ability", or "KSA" for short. The view has been defended and developed by Ernest Sosa, John Greco, Alan Millar, and to certain extend by Duncan Pritchard, and Julien Dutant. I will focus on Sosa's version here, which I consider a species of

mild anti-conjunctivism, because it follows Zagzebski's strategy of blocking the Gettier-type examples. It does so without pursuing such a close parallel between the epistemic domain and the domain of morals.

In his recent work, Sosa (2007, pp. 22-44) appreciates the danger of the Getter cases, and possibly of the primeness argument, and concentrates upon a more restrictive connection between the two components which would suffice for knowledge. First, Sosa takes it that believing is a performance like other skilful performances, such as archery. Then, he suggests that our beliefs are performances that have a so called "AAA structure". Sosa claims that each performance has a goal, and can be assessed according to three main criteria: *accuracy*, *adroitness*, and *aptness* (hereby the AAA structure). Accuracy is a criterion of successfully achieving the goal, adroitness is a requirement of having the ability for successfully achieving the goal, and finally, and crucially, aptness is achieving the goal as a result of applying the ability under question. More precisely, the relation of "aptness" is a causal relation between competence and reaching truth where competence is a cause and reaching truth is an effect. To clarify the notion of aptness, Sosa draws an analogy between a cognizer and an archer. When an archer hits the target as a result of her skill, her shot is apt. When a cognizer reaches a true belief as a result of her competence, her belief is apt, and therefore amounts to knowledge. Apt belief, according to Sosa, constitutes the necessary and sufficient condition for knowledge. Animal knowledge is now defined as apt belief, and reflective knowledge is subsumed under a double-aptness criterion: it is an aptly formed meta-belief about the first level apt belief. Now Sosa's strategy against the Gettier problem is simple. Gettier cases are explained as cases where one's belief is adroit and accurate, but it is not accurate *because* adroit, i.e. not apt. The same explanation can be given as to why a subject of deviant causal chain agent does not know. Sosa identifies adroitness with intellectual virtue. Unlike Zagzebski, though, Sosa thinks that an act produced by intellectual virtue is not always

perfect in all respects. One can act out of virtue, and yet fail to have knowledge, like in the Gettier cases. In this sense, the Gettier cases are not parasitic on this account of knowledge; they are needed to reveal the importance of the causal determination of truth.

John Greco proposes us to understand success from ability in terms of explanatory salience. According to Greco (2009), attributions of knowledge through ability are dependent on the practical reasoning context – taking into account practical considerations within the context. For example, a sensitive attributor would judge that:

> Derek Jeter has the ability to hit fastballs relative to normal environments for playing baseball. He does not have the ability relative to an active war zone, where he would be too distracted to focus on the ball. (p. 21)

An attributor is to grant knowledge not just when the subject has an ability but also iff the case is such that the ability of the subject is responsible for her cognitive success, whereby Greco understands the causal contribution to success in explanatory terms: *S* has knowledge in cases where her believing *p* by exercising an ability explains in a salient way why *S* has a true belief. Greco accounts for 'a salient explanation' partially as a function of our interests and purposes as information-sharing beings. He also mentions that for explanatory salience the exercise of ability should be partially under the subject's control, and in her power to improve upon it.[15]

It is certainly true that contextual considerations are to be taken into account when we judge cases of success from ability. However, Greco's purpose here is to explain knowledge, not to provide a definition of knowledge. Unlike that, we are pursuing a project of defining knowledge. The notion of explanatory salience is not sufficiently strong to play the crucial role in a definition of knowledge. This is so, because explanatory salience does not grant success from ability. To take a radical example, the subject of attribution might be a zombie acting in a given context as if she adequately achieves truth from her abilities. The attributer

[15]A theory of explanatory salience is developed by Van Fraassen (1980) pp. 97-158

would be justified in believing that this individual achieves knowledge as a result of her abilities. But she would be wrong.

2.5 Primeness of apt belief

Finally, let me demonstrate the compatibility of the aptness account of knowledge with Williamson's primeness argument. I will try to show that the relation of aptness entails primeness since it is not a mere conjunction of external and internal, but entails the proper idea of collaboration between the two.

Here is my argument that aptness entails primeness, constructed on the basis of Williamson's original argument.Consider three different cases: α, β, and γ. Suppose that there is a target T in front of a subject S, and there are no external obstacles to hitting the target (no wind and the like). The shooter is aiming at it with two guns in her hands. One of the guns does not function. So, in case α: *S*'s brain is properly sending a signal to the right hand holding a working gun. *S*'s brain center of sending a signal to the other hand is impaired. However, she has a muscle twitch of that hand which makes her pull the trigger. As it happens, the gun in her left hand does not work. As a result, *S* hits the target aptly. In case β: *S*'s brain is properly sending a signal to the left hand holding a working gun. *S*'s brain center of sending a signal to the other hand is impaired. However, she has a muscle twitch of that hand which makes her pull the trigger. As it happens, the gun in her right hand does not work. As a result, *S* hits the target aptly. Now, case γ is internally like α, and externally like β. In γ: *S*'s brain is properly sending a signal to the right hand holding a non-working gun. *S*'s brain center of sending a signal to the other hand is impaired. However, she happens to hold the working gun in that left hand. By accident, a muscle twitch makes her pull the trigger of the gun. As it happens, the gun in her left hand does not work. She hits the target, but her hitting

the target in inapt. This generalizes to other abilities. Therefore, the relation of aptness is not vulnerable to the primeness argument.

It is more difficult to establish that the definition of knowledge as apt belief is not circular. However, we have a reason to believe that it is not, at least because it is illuminating. On the one hand, adroitness is a natural tool for reaching truth, on the other - it is fallible and hence it does not presuppose knowledge. Also, a deviant causal chain cannot be generated without violating the condition of aptness. Now, if the last consideration gives us reason to think that the definition is circular, this suggests that every definition that is immune to counterexamples, and therefore is a good definition is circular. If this is a reason to avoid definitions, nothing can help against it.

CHAPTER 3: RYLE'S TURN AND THE GENERALITY OF THE APTNESS ACCOUNT OF KNOWLEDGE

> Marksmanship is a complex of skills, and the question whether he hit the bull's eye by luck or from good marksmanship is the question whether or not he has the skills, and if he has, whether he used them by making his shot with care, self-control, attention to the conditions and thought of his instructions.
>
> Gilbert Ryle (1949, p. 45)

> (I)n speaking of a justified belief we are saying something rather like "Good shot!" which someone might sincerely and correctly say despite being opposed to gun possession and to shooting.
>
> Ernest Sosa (2007, pp. 66-67)

3.1 Introduction

As we saw in chapter 1, debates about knowledge are centered on issues of justification of our beliefs in various propositions. The study of skills responsible for arriving at such beliefs, and in some cases obtaining knowledge, was often discussed in the descriptions of the so-called context of discovery, taken to be secondary in relation to the context of justification. Reliabilism, in contrast, gives the issue of how we arrive at our beliefs some importance. But even generic reliabilists do not take it as a primary task to analyze the structure of reliable processes themselves. In this respect, virtue epistemologists are the first to bring forward the idea that reliable abilities (in particular cognitive faculties) have normative properties represented by their structure. For instance, Sosa (2007, p. 113) takes it that our competences have the following normative components: constitution: e.g. the perceptual skill (the seat of the competence); condition: e.g. being awake and sober; and situation: e.g. enough light. Now, if epistemic normativity is a status by having which a true

belief constitutes knowledge, then these three components enter into the constitution of something, which according to Sosa (2007, p. 88) has "fundamental epistemic worth: namely, apt belief, true because competent". The study of skills from this perspective would be analogous to the study of the structure of justification in traditional epistemology.

This chapter is supposed to make a modest contribution to the work in that direction. First, my claim is that justification is not separate from skills and capacities, and second that skills and capacities constitute the essential basis, that is both temporarily and normatively prior to the traditionally appreciated knowledge-that. To give an example, knowledge that *Köszönöm* means *Thank you* in Hungarian is grounded on, or constituted by mastery of language by the native speaker, and this in turn is due to a complex of abilities of speaking and understanding.

I will explore the notion of "intelligent ability" instead of intellectual virtue since virtues are features of one's cognitive character, or excellences that are a more restricted source of knowledge. What I want to do first is to deflate Sosa's criterion of adroitness. Indeed my arguments are compatible with the notion of virtue as function, but I would like to claim that we can develop various abilities beyond the proper functions of our faculties. And these abilities can serve as sources of knowledge. For instance, one's ability to dance tango is also a source of one's knowledge how to dance tango. But dancing tango is not a proper function of a faculty. The criterion of aptness, as originally meant by Sosa, requires reaching success from Platonic virtue. However, we can define a sense of apt causal determination at various levels where success is caused by a quality of a cognitive state of x which is not a virtue. For instance, in thermometer, a mechanism causes successful display of temperature. Call this apt** causal determination. Apt** causal determination is obviously too weak for knowledge, and it is even too weak to be a cognitive state of a lower type. Nevertheless, there are various possibilities between virtue-apt and apt** types of causal determination. For

instance, we have a relation in which a cognitive mechanism which is not a virtue brings a fact to someone's mind. Call this apt * causal determination. I want to claim that apt*class of causal determinations is pretty large, and that this class is responsible for the core of our knowledge. It represents success from intelligent abilities, which is a minimal condition for knowledge. This class does not coincide with the virtue-apt class of causal determinations though, because apt* causal determinations include more deficient kinds of success, such as success caused by an improving ability; or success caused by an intelligent ability which does not display the excellence needed for virtue (like the calculating ability of someone who is not as good in math as an excellent mathematician). I suppose that virtue even taken as a proper function entails functioning well. Apt* is meant to cover also cases where one occasionally forms true beliefs on the basis of a mechanism that properly brings them to mind, but the mechanism is unstable in doing so. Sosa's requirement of aptness seems to cover such cases too:

> If the act is due to a competence exercised in its appropriate conditions, its success may be due to luck in various ways. It may be just an accident that the agent retains his relevant competence, for example, or that the conditions remain appropriate. Either way, the act fails to be safely successful, since it might too easily have failed, through lack of the required competence or conditions. It might still be apt, nevertheless, indeed attributably, creditably apt. (Sosa 2007, p. 81)

The general factor behind apt* determinations is just intelligent ability. More precisely I would like to suggest that in his discussion of knowledge how Gilbert Ryle offers a good explanation of what an intelligent ability which is here claimed to crucially participate in apt * type of causal determination has to be.

Now, the second major point of this chapter is to draw some consequences about the generality of the aptness account of knowledge from Ryle's debate of knowledge how. In particular, I will argue that the tendency of epistemologists to focus narrowly upon knowledge that is misguided. Let me call this narrow focus "intellectualism". The aptness* account of

knowledge offers a potential for enlarging this narrow focus and to gain territory, so to speak, because the principle of aptness* is an adequate requirement for our practical knowledge too. Sosa himself endorses this line of thought (in Bonjour and Sosa 2003 pp. 99-102).

By extending truth to the more general species of success, and intellectual virtue - to the general category of intelligent ability, we can obtain from Sosa's aptness account a generalized aptness account that covers also instances of knowledge how:

A*: *S* knows that *p*/how to *ψ iff S* grasp the truth of *p*/successfully *ψs* as a result of *S*'s intelligent ability.

Since truth is a subspecies of the general notion of cognitive success, by means of the KSA principle we can obtain a fully general definition of knowledge:

KSA*: Knowledge is success from intelligent ability.

3.1.1 A reason for generalization

The kind of generalization I am proposing would be needed only if we find compelling the claim that at least some instances of knowledge are non-propositional. In this case, having a fully general definition of knowledge would be desirable.

Notoriously, in his *The Concept of Mind*, Gilbert Ryle suggested that knowledge is not exhaustively propositional, and that there is another species of knowledge, knowledge-how. A *prima facie* reason to accept Ryle's distinction is that, at least sometimes, we use the phrase 'know how' to refer to practical knowledge, such as: "I know how to swim", "I know how to juggle" which is arguably not reducible to propositional knowledge. One way in which it is not reducible is that while for claims of the sort "I know that swimming is an Olympic sport" one can give a convincing justification in the form of other propositions or explanation, in some contexts at least it is not enough to provide support for claims of the sort "I know how to juggle" if one offers just an explanation of how one is supposed to juggle. In such contexts

what is required is either that one has to show you the trick, or you have to check how it works in practice. The difference in the two kinds of justifiers is established by the fact that transition from knowing the way in a propositional sense to knowing it in a practical way is not automatic. The first does not logically entail the second. I may know how one is supposed to juggle without knowing how to juggle myself. And vice versa: I may be able to juggle myself, but not in a position to explain you how one is supposed to do it. So, at least in some contexts the ultimate justifier for knowledge how is of purely practical nature. This gives us a *prima facie* reason to accept the claim that there is a kind of knowledge which does not originate in propositional knowledge and is not itself propositional.

If Ryle is right in spotting such kind of knowledge, this kind of knowledge is a trouble for the intellectualist who wants to give an account of knowledge exclusively in terms of consideration of propositions. The radical way to handle the unwanted residue of knowledge from an intellectualist perspective is to argue that knowledge how is a species of knowledge that. Jason Stanley and Timothy Williamson (2001) present a technically lucid defence of this claim. More recently, John Bengson and Marc Moffett (2007) continue the line. The more modest way to handle the residue from that perspective is to throw it out of the board of epistemology, and to claim that epistemology should be concerned with an account of propositional knowledge only. This is a concession which we want to avoid, because the residue is not an insignificant part of our knowledge. Our practical knowledge is a substantial part of what we know, and it is, as Ryle further suggests in his negative argument against the intellectualist doctrine, the source of the value of our theoretical knowledge.

I propose not to dismiss the residue, but to use it in favour of a more radical ability based approach to knowledge inspired by Sosa's and Greco's virtue epistemology. The step that I am going to propose is radically anti-intellectualist. I argue that all knowledge is based on the exercise of intelligent abilities. So, my point is not to stress the distinction between

knowledge how and knowledge that, as Ryle's point is usually interpreted. On the contrary, I straightforwardly claim that any kind of knowledge is based on intelligent abilities understood pretty much in the way that Ryle described knowledge how.

3.2 Preliminaries

I should make two preliminary remarks. First, Ryle's theory of mind has been largely criticized for its ubiquitous anti-mentalist bias. And while his book *The Concept of Mind* (1949) is a source of some important ideas rediscovered, and defended by contemporary epistemologists, this fact is rarely paid attention to, because Ryle himself used these ideas as premises for a doctrine of a non-compelling character, which looks unattractive to many. In the chapter "Knowing How and Knowing That" (pp. 25-62), Ryle combines two famous agendas: the defense of the importance and irreducibility of knowledge-how and knowledge-that, and a kind of Wittgensteinean anti-mentalism (rejection of the "ghost in the machine"). I would like to resurrect only the first agenda. I shall steer clear from the second (behaviorist) agenda. I think that because of the pejorative behaviorist interpretation of his whole work (which is correct, but it seems to me too restricted), Ryle's contribution to contemporary epistemology has been underestimated.

Secondly, I am strictly interested in Ryle's idea that there is a kind of knowledge which is a direct result of the exercise of intelligent abilities, independently of whether it is properly called knowledge how or knowledge that. I think this still captures partly Ryle's original idea. What is required from knowledge-producing abilities is that they are intelligent abilities, unlike the ability to breath or digest; and that they are refined in regard to doing their jobs. The claim is that any kind of knowledge is a direct result of exercise of intelligent abilities. This theory plays an important role in my last two chapters.

3.3 Ryle's turn

I take it that my generalization project originates in a tradition which started with Gilbert Ryle, or presumably before that, and it is carried on by contemporary virtue epistemologists. The tradition I am talking about is the tradition of seeing propositional knowledge as based on possession of intelligent qualities or competences. In the spirit of that tradition, virtue epistemologists indeed advise us to undertake an anti-intellectualist move, i.e. to take propositional knowledge as based on intelligent virtues (qualities) of the cognizer, and not on classical type of justification of *p* by means of other propositions. So, the tradition shared by Ryle and contemporary virtue epistemologists I am referring to is associated with the idea that intelligent qualities of the cognizer rather than justificatory status of beliefs or propositions constitute the ground of propositional knowledge. Starting with Descartes' *Regulae ad directionem ingenii* (1619-1628, in1998), passing through Lewis Carroll's famous "What the Achilles Said to The Tortoise" (1895), the idea was given flesh by Gilbert Ryle. He turned our attention to a potentially significant change in the focus and the scope of epistemological investigation, and thus anticipated the shift of the focus from the epistemic status of our beliefs to the epistemic status of agents implemented by contemporary virtue epistemologists. Unlike Ryle though, contemporary virtue epistemologists invite us to undertake the suggested move without discarding the rich problematic which is already on epistemology's offer.

In his chapter "Knowing How and Knowing That" in Ryle (1949) develops a theory of intelligent ability. It is grounded in his distinction between knowledge how and knowledge that, which is the first systematic attempt to do the generalizing step I have been talking about. Let us see how.

Looking at Ryle's own terminology can be helpful as an initial glance to the distinction. The term "intellectual", as Ryle uses it, relates to the activity of "considering

propositions", and is opposed to "intelligent" meaning a display of intelligent qualities, or learnt abilities. In particular, Ryle thinks that *knowledge that* is knowledge of truths, while *knowledge how* is intelligent ability (p. 28). So generally, the distinction between knowledge that and knowledge how is a distinction between the domain of intelligence, and the apparently narrower intellectual domain. The distinction is motivated by a simple common-sense idea:

> Intelligent practice is not a step-child of theory. On the contrary theorizing is one practice amongst others and is itself intelligently or stupidly conducted.

and further:

> it is important to correct from the start the intellectualist doctrine which tries to define intelligence in terms of apprehension of truths, instead of apprehension of truths in terms of intelligence. (p. 27)

To clarify, the intellectualist doctrine, according to Ryle, is the view that consideration of propositions is the kernel of intelligent action. Considering propositions is the specific activity which yields to knowledge that. More precisely, knowledge that is a relation between a thinker and a true proposition. Traditional epistemology, which conceives of justification as justification of propositions by means of other propositions is closely associated to the intellectualist doctrine, because it treats knowledge as restricted to knowledge of truths.[16] "Ryle's epistemological turn", as I call it, is a rejection of the basic order of priority in traditional epistemology. Ryle argues that consideration of propositions is not the kernel of intelligent action, but rather intelligent action underlies the epistemically valuable consideration of propositions.

[16] According to a broad epistemological interpretation of knowledge that, it has two specific characteristics: i) it requires acceptance of truth or falsity of a proposition, ii) it requires sufficient reason for accepting the proposition (widely interpreted as justification).

My particular enemy here is the traditional epistemologist as described above rather than the proponent of the intellectualist doctrine as originally meant by Ryle.[17] There is a specific kind of regress which can be generated if one embraces the view that a proposition can be supported only by means of other propositions. Imagine that I claim that Mary is in love with *X*, a friend of yours. You would probably be interested how do I know that? I tell you that I saw them talking in the corridor. Call my statement "*L*". Since this does not seem to be a sufficient reason in support of my statement, you ask me what makes me think that talking in the corridor with *X* is a good evidence for being in love with *X*? Then I reply that Marry was smiling at *X* while talking. Call this "*R*". You feel dissatisfied with my answer, and you want to know what makes me think that smiling was an expression of love as opposed to being in a hilarious mood? Then I reply that she was smiling for too long. Call this "*E1*". You still might wonder how smiling for too long makes any difference, so you ask for further evidence about the probabilistic relation between *E1* and *E*. So, in order for my claim (*L*) to amount to knowledge, I must have a reason to believe it *R*. But I must also be justified that *R* is a reason for *L* on the basis of a further reason *R1* which explains why *R* makes *L* probable, and so on ad infinitum. Hence, a regress of reasons is being generated. In order to terminate the regress, classical foundationalists assume that in order for *S* to know that *p* there must be some non-inferential reason that terminates the regress of justification.

Bonjour famously claimed that this not a good solution to the regress problem, as I already mentioned that in chapter 1. Here is why. Foundationlists assume that a belief is non-inferentially justified if it has some characteristic *X*. BonJour (1985) then argues that the mere fact that the belief has *X* could not, even in principle, justify the believer in holding the belief. The believer would also needs to have access to the fact that the belief in question has *X* and also some reason to hold that beliefs of this sort are likely to be true. At least one of these

[17] Ryle thought that the proponent of the intellectualist doctrine is committed to the claim that in order to know that p, a cognizer has to *contemplate* propositions of endlessly complex kind. I do not take my enemy to be

propositions could only be known through inference, and thus the regress of reasons makes its way through. Is there another way out of the regress? Sosa suggests that groundedness in virtue can terminate the regress. I would like to suggest that an intelligent act can terminate it. What makes a reason a good reason is that it is a result of intelligent act.[18] I think this was also what Ryle wanted to establish, although he reached that conclusion through a different kind of regress.

3.3.1 Introducing intelligent abilities

An intelligent act that yields knowledge results from an intelligent ability. So, the next question is what is an intelligent ability? What we want to know more precisely is what an intelligent ability behind the apt* causal determination should amount to. I suggest that we look at Ryle's positive account of knowledge how (pp. 40-50), where he presents a reasonable theory of intelligent ability. Also, he suggests that something like knowledge how is the fundamental basis of our knowledge. According to Ryle's positive theory, knowledge how is a complex ability to apply rules of successful performance in practice, and which is a product of practice. In most cases, Ryle suggests, knowledge-how includes a set of sub-abilities of different kinds: not just physical, but also mental, such as memorizing, thinking, and reflecting over one's own experience. For instance, "knowledge of how to tie a clove-hitch includes abilities not just of tying clove-hitches and in correcting one's mistakes, but also in imagining tying them correctly, in instructing pupils, in criticising the incorrect or clumsy movements and applauding the correct movements that they make, in inferring from a faulty result to the error which produced it, in predicting the outcomes of observed lapses, and so on indefinitely." (p. 55). Thus, Ryle sees knowledge how as a complex multiple-track diposition

committed to this assumption.

[18] In terms of my general argument, I would like to clarify that intelligent act is sufficient to stop the regress of justification. It does not need to be an act of intelligent ability.

that has no (easily statable) necessary or sufficient conditions. It is a disposition to react diversely depending on the circumstances. He contrasts such intelligent multiple track dispositions with "single track dispositions like a reflex or a habit" (p. 44). Single track disposition is a disposition to perform an act (automatically) in order to fulfill a task, such as pressing a button. The exercise of a single track disposition does not require thinking or alertness.

As a matter of fact some habitual performances can result from multiple track dispositions. Take, for instance, driving a car routinely. This performance is certainly an expression of knowledge-how, but I am not sure how this fact squares with Ryle's theory. One way to accommodate it is to say that habitual performances from multiple track dispositions are expressions of knowledge because they inherit the previously invested propensities. I think that it is essential to appreciate that in this way propensities are also embedded in the routine exercise of all intelligent abilities. All intelligent performances are intelligent because either they are exercised with propensity, or because they inherit invested propensities (in case no particular alertness is required for the concrete action). For instance, skilful driving entails a routine execution of complicated movements when circumstances are normal.

Ryle's theory of knowledge how is based on his conception of the nature of intelligent act. Intelligent act, according to Ryle, is an act done with "presence of mind" or with paying heed to the relevant factors in a context such as attention to what one is doing, or readiness to take into account a relevant change in the circumstances during a cognitive performance. Each such act is performed with a special "*modus operandi* " or "propensity", like care, self-control, attention, noticing, taking care, applying one's mind, concentrating, putting one's heart into something, thinking what one is doing, alertness, interest, intentness, studying, trying, etc. Propensities do not just accompany intelligent acts, but are a tool for achieving

knowledge how. What enables the formation of and sustainment of multiple-track dispositions is the sequence of such intelligent acts in certain domain. Now, if consideration of propositions derives its epistemic value from the intelligence of the act of considering, then intelligent acts of considering propositions are our tool of acquiring propositional knowledge too.

There are opponents to the theory that knowledge how is different from knowledge that. A well known defense of intellectualism along these lines is offered by Jason Stanley and Timothy Williamson (2001). The two authors suggest that knowledge how is knowing a Russellian proposition ("that…"), under a practical mode of presentation ("how…"):

x knows how to ψ if and only if for some way *w* of ψ-ing

i. *x* knows *w*, and

ii. *x* knows that *w* is a way of ψ-ing under a practical mode of presentation.

Thus, the two authors contend that knowing how always reduces to propositional knowledge, as documented by the quote:

> it is simply a feature of certain kinds of propositional knowledge that possession of it is related in complex ways to dispositional states. Recognizing this fact eliminates the need to postulate a distinctive kind of nonpropositional knowledge.(p. 430).

I would like to point out that this kind of "reduction" still leaves the crucial issue open. The reducing element consists of a Russelian proposition and its practical mode of presentation. The Russelian proposition is not entertained as such, but through a mode. So, the fact that it is a proposition does not tell much about its mental "presence". What is left and what is crucial for knowledge how is the practical mode. So, the crucial issue returns, so to speak, through the backdoor: is the practical mode fundamental in relation to the non-practical? In the cases of physical dexterity, e.g. skilful bicycle riding it is quite obvious that the answer is "yes". In the case of cognitive capacities one can argue for the analogous

conclusion from the well-known threats of circularity or regress familiar from Ryle's original discussion.

Next, the whole work in Stanley and Williamson's reduction argument is done by the practical mode of presentation and the allegedly propositional part is played by such a thin proposition like a Russellian proposition that it all amounts at the end of the day to a practical mode of a presentation of a fact, which does not really prove that knowledge that is to be taken as fundamental. Evidence that the same know how is exercised in imagining a movement as in performing it is shown by the fact that rehearsing movements in imagination improves physical performance.

Moreover, Stanley and Williamson's practical mode of presentation is essentially linguistic: "In what follows, we shall speak of modes of presentation as being associated with certain linguistic constructions…" (pp. 427, 428). If modes of presentation are linguistic constructions, they essentially involve concepts. But, it is hardly sustainable that we have concepts for practical dispositions. Concepts, at least taken in classical sense, are theoretical, not practical entities. But if modes of presentation do not involve concepts, then how can they involve propositions?

Finally, Ian Rumfitt (2003) shows that the argument of Stanley and Willamson is driven by intuitions of an English language speaker, and that speakers of other languages do not share the same intuitions. French, like many Romance languages, naturally expresses knowing how via an infinite complement to the 'know' verb, as in 'Il sait nager'. And in Russian 'know how' is translated by a verb that is not identical to the propositional knowledge verb used in Russian. So, it seems that eventually Stanley and Willamson's argument against Ryle's claim is unsound.

3.4 Virtue epistemology as a part of Ryle's tradition

One of the main assumptions of virtue epistemology is that normative properties of agents, i.e. agent's virtues, and not normative properties of an agent's beliefs are the primary source of epistemic value. This is how virtue epistemologists claim to differ from the preceding tradition in epistemology. We can see this claim as congenial to Ryle's turn just described, namely of his claim that there are no proper criteria for intellectual success independent of its intelligent basis. This section aims at introducing more common points between Ryle's turn and Sosa's and Greco's virtue epistemology.

I will start by quoting a table from Sosa's book (2007, p. 78*),* where Sosa draws an analogy between epistemology and a theory of practical skills:

EPISTEMOLOGY	ARCHERY CRITICISM
Beliefs.	Shots.
Concepts.	Bows, arrows.
Believers and their ways of forming and sustaining beliefs.	Archers and their ways of shooting.
Schools, methods of teaching, modes of inquiry. Epistemic communities, criteria for rating and promoting.	Communities that preserve, supplement, and transmit the lore of archery, and that honor accomplishment in accordance with criteria.
True beliefs.	Accurate shots.
True beliefs might vary in epistemic respects; for example, some have more content than others, being more specific.	Accurate shots might vary in respects relevant to archery; for example, some come closer to the bull's-eye than others.

In this table, Sosa clearly commits himself to the view that beliefs are a kind of performances. At many places in the same book, Sosa identifies beliefs and performances. For instance: "Beliefs too might thus count as performances, long-sustained ones, with no more conscious or intentional an aim than that of a heartbeat...Beliefs fall under the AAA structure,

as do performances generally." (p. 13), "This second lecture proposes a virtue epistemology that distinguishes between the aptness and the safety of performances generally, and of beliefs in particular. " (p.23) "A virtuous performance, whether a correct belief due to intellectual virtue or a right action due to practical virtue, will involve both the agent's constitution and his situation. " (p.45), "I am here opting for something like the ballerina's graceful pas as correlate of the knower's believing. The apt believing is a performance-immanent value, unlike the cup of coffee vis-`a-vis the doings of the barista." (fn. 48 p. 48).

Propositional knowledge is presented by Sosa as a kind of performance, subordinated to criteria and norms which hold for any intelligent performance. Also, notice that the archery is a typical example of knowledge how. So, the table suggests that epistemology and practical criticism share the same kind of normativity.

Knowledge entails the presence of corresponding intelligent ability. Suppose I formulate in the course of ordinary conversation a sentence *S* in my mother tongue in this case it is Bulgarian, and somebody asks me in which language it is. I know immediately that it is a sentence of my mother tongue. My judgment to this effect is accurate, because adroit. I have the ability to produce and recognize sentences of my mother tongue like for instance S. In this case the ability is often described in terms of knowing-how. I know how to produce and recognize the sentences in question. This is a cognitive analogue of more purely practical knowledge-how, for example one's ability to dance tango. Similarly, the ability of my mathematician friend to arrive at true mathematical beliefs is standardly described in terms of knowledge-how: he knows how to prove a theorem, for instance. Obviously, a lot of ordinary cognitive abilities belong to this genus of knowledge-how. If Sosa is on the right track, knowledge-that which they produce (my knowledge that the sentence belongs to my mother tongue and my friend's knowledge that something is a theorem of mathematics) is knowledge because the knowledge-how essential for their production is itself adroit.

Also, Sosa's competences and Ryle's knowledge how are structurally analogous. As I have already mentioned, Sosa thinks that competences are made of the following three components: constitution: the archer's skill (the seat of his competence); Condition: being awake and sober; Situation: enough light, normal wind. Finally, there is always context that settles the aim. Likewise, Ryle thinks that knowledge how is constituted by intelligent ability which consists of: the skill (presumably the seat of the ability), propensities with which the skill should be exercised (a more fine-grained version of Sosa's "condition"), in the presence of normal circumstances. Of course, some practical goal is also supposed to be settled in the context in which the ability is being exercised.

We have at least four more reasons to think that Ryle's turn is congenial to virtue epistemology as promoted by Sosa and Greco. I outline them briefly in the following paragraphs.

Anticipating reliabilism. As we saw, Ryle's turn amounts to the idea that intelligent properties of the agent ground knowledge of truths. The idea that intelligent properties (i.e. reliable processes, or faculties) as opposed to internalist justification of our beliefs ground knowledge of truths came to be known later on as 'reliabilism'. According to process reliabilism, what justifies *p* is *not* the consideration of any proposition, but a *quality* of a process. Thus, Ryle had opened a way to what was going on later on in epistemology.

O*vercoming generic reliabilism*. Remember that by claiming that knowledge is grounded in faculties, rather than processes Sosa and Greco claim that their theory overcomes the drawbacks of generic reliabilism, revealed by the objections such as Mr. Truetemp (Lehrer 1990), The Brain-Lesion case (Plantinga 1993), and the Clairvoyant case (BonJour, 1985). Ryle's original idea is inspired, or at least goes together with the notion of intelligence as *ability or disposition*, rather than with the notion of a reliable *process*. Ryle's theory of intelligent ability is equally armed to overcome the same objections. Ryle's intelligent ability

is distinguished from physiological processes such as blood circulation or breathing, or mechanical processes such as showing the temperature, and even from some well-drilled animal performances. Let me briefly address Mr. Truetemp, and the clairvoyant objection.

Take first Lehrer's (1990) objection against generic reliabilism, known as Mr. Truetemp. The story concerns Mr. Truetemp who has no idea that a thermometer has been implanted in his head. As a result of that he has regular reliable thought about the ambient temperature, and he does not know where they come from. The generic reliabilist is bound to say that Mr. Truetemp has knowledge as to what the temperature is. But, this is commonly taken to be implausible. On Ryle's theory of intelligent ability, the thermometer's ability to show the temperature is not an intelligent ability, it is like physiological processes. It cannot be displayed with propensities. So, the mere fact that Mr. Truetemp has a reliable thermometer in his head is not sufficient for his knowledge of the temperature. He has to exercise that ability with certain propensities.

In its most general version, the clairvoyant objection against generic reliabilism is the objection that if *S* obtains a clairvoyant ability out of the blue, and start forming reliably true clairvoyant beliefs, these beliefs will not amount to knowledge, at least initially. Ryle's reply to the clairvoyant objection could be that the initial clairvoyant ability can develop into an intelligent ability, if it is exercised with the necessary propensities, for instance caution, and sensitivity to counterevidence, etc.

Another common point shared by Ryle and Sosa is *the idea that execution of a criterion is primary to considering the criterion.* The common idea is that the display of intelligent abilities presupposes application of norms which are most often unknown to the cognizer. On his side, Ryle describes knowing how as knowing rules (right procedures) in an executive way: "the application of the criterion of appropriateness does not entail the occurrence of a process of considering this criterion." (p. 31). He thinks that "actions done

from motives can still be naive, in the sense that the agent has not coupled, and perhaps cannot couple, his action with a secondary operation of telling himself or the company what he is doing, or why he is doing it." (p. 111). On the other side, Sosa (2007) describes animal knowledge as being motivated by reasons which one could not fully appreciate:

> The archer is motivated by reasons to release the arrow when the bow and arrow are held just so. He may be unable to articulate these reasons, but we cannot plausibly require that our reasons must always be articulable, lest we deprive ourselves of reasons that matter to us as much as anything. We identify a loved one, for just one example, in ways that we could not articulate fully. (p.84)

Animal propositional knowledge obeys the same principle by analogy. The right kind of justification, according to Sosa, is such that it goes through application of norms (virtues) which are first subconscious and then brought out by reflection. These norms are characteristics of our intellectual virtues, responsible for the production of knowledge.

3.5 Personal and sub-personal abilities

Finally, I want to discuss a last common point between the two authors, and to say what I find problematic in it. Both Sosa and Ryle are committed to the claim that propensities/virtues determining intelligence are the agent's characteristics. Ryle's propensities determining intelligence are to be seen as characteristics driving the epistemic performance on the agent's level. They are determined by the "agent's frame of mind": in performing an intelligent act one has to think what one is doing. (p. 29) Sosa's and Greco's virtues, on the other hand, are claimed to be reliable faculties creditable to the agent.

However, it seems to me that a significant number of propensities involved in our spontaneous intelligent actions, and in the formation of intelligent skills are such that the agent is unaware of them. We have naturalistic reasons to think so. When it comes to the formation of intelligent skill, the brain exclusively does the job in building up a neural network. The building up of neural network is a process which is only indirectly and only

sometimes influenced by agent's conscious effort. We should not think that propensities are always intentionally invested by the agent in the formation of her skills. If I watch how my flat mate prepares the same cake every week, I can unintentionally learn to prepare it myself.

We can eventually break down the issue of the independence of intelligent abilities into two sub-issues: abilities on the conscious "personal" level; and abilities on the unconscious, "sub-personal" level (on that level we usually do not talk in terms of knowledge how and knowledge that, but of their sub-personal correlates). Compelled by his anti-mentalist agenda, Ryle considers mainly the first sub-issue, but it is easy to extend the consideration to the second one as has been done abundantly in the literature.[19]

Intelligent skill of Rylean type at the conscious "personal" level is relatively easy to illustrate and defend. When I formulate in the course of ordinary conversation a sentence in my mother tongue, I am not aware of applying consciously any rules of grammar let alone of mobilizing knowledge that such rules hold for my language. When I conclude from the truth of a particular conditional and the truth of its antecedent to the truth of its consequent I am normally not explicitly aware of applying modus ponens, i.e. detachment rule. My reasoning is "blind" (Boghossian and Williamson 2003). Let me call such knowledge "personal-level knowledge how".

At the unconscious "sub-personal" level different processes can accompany and sustain mental activity that corresponds to the personal level intelligent skill. For instance, Chomsky has famously (but only at some places in his work) argued that formulating a sentence is actually sustained by unconscious process having access to the explicitly represented rules of language. His opponents claim that such rules are only implemented in our cognitive system without being represented. At this level the debate between, so to speak, intellectualist and the anti-intellectualist reiterate. At the unconscious sub-personal level one

[19] For an overview of the debate, and the author's position see Devitt (2006)

can have either guidance by explicitly *represented* rules, a kind of victory of knowledge-that at the sub-personal level, or guidance by merely *implemented* rules.

This table represents how propositional and practical knowledge is placed on the two levels:

	Knowledge-that	Knowledge-how
Personal	T1: conscious propositional knowledge	H1: conscious practical knowledge
Sub-personal	T2: guidance by explicitly *represented* rules e.g. classical symbol systems	H2: guidance by merely *implemented* rules e.g. connectionist systems

So, I will take "intelligent ability" described in the above sense to cover both skills grounded exclusively in unconscious sub-personal level, and skills that demand sophistication at the conscious personal level. Usually, the later are taken to evoke the idea of "knowledge-how". However, seeing, hearing, or touching seem to be sub-personally grounded, but they require equally sophisticated abilities as the ones that stand for Rylean kind of knowledge how. One reason for that is because perceptual knowledge is based on discriminative and recognitional skills that evoke the idea of knowledge how. An ability to discriminate between cypress trees and spruce trees entail know-how, as well as the ability to recognize each of these species of trees. We can say: 'Peter knows how to recognize spruce trees.' Central vision routinely exercise discriminating and recognizing. On the other hand, a lot of non-perceptual abilities tied to inferences and even to simple insights (e.g. my simple linguistic insight) are routinely characterized as instances of knowledge-how too. This fact eventually can help us to see our faculties on par with other practical abilities, intentionally sophisticated through practice and representative of knowledge how. But indeed one's basic faculties are rarely intentionally sophisticated by the agent, they are sub-personally so sophisticated, as well as many of our practical skills are.

In conclusion, I suggested that if epistemologists want to avoid regress of justification, they have to embrace a kind of foundationalism about intelligent qualities of our epistemic performances. This line of thought was anticipated by Ryle, and can be seen as initiating a tradition which was slowly taken up by contemporary epistemologists: first generic reliabilists, and then, more significantly, by virtue reliabilists. On the other hand, Ryle's own account of intelligent ability were restricted to explaining the nature of knowledge how, but indeed it can be seen as providing a more general explanation of our intelligent abilities: not just practical, but also theoretical. The difference between practical and theoretical knowledge from this perspective is in the class of intelligent abilities involved in the production of each. Hence, we can use Ryle's model of intelligent ability to describe the kind of adroitness* involved in apt* causal determinations. Ryle's theory of intelligent abilities can be seen as more down-to-earth and generalized alternative of intellectual virtues, because it does not presuppose the intervention of one's cognitive character into the performances of singular abilities. Yet, it gives us a minimal subjective element over and above the mere reliability in the form of propensities. Next, such a theory of adroitness* has at least equal potential of solving some of the perennial epistemological problems as the virtue account does. I will try to demonstrate this in the next two chapters. Finally, the notion of intelligent ability allows us to ascribe basic credit for knowledge to one's sub-personal cognitive system, and not to ascribe the whole credit for knowledge undeservedly to cognitive agents.

CHAPTER 4: EPISTEMIC LUCK AND SAFETY OF KNOWLEDGE

> Empirical knowledge is possible only because the environments in which we try to acquire it exhibit uniformities that we take for granted.
>
> Alan Millar
> (in Pritchard, Millar, and Haddock 2010, p.186)

4.1 Introduction. Achievement and luck

It is possible for a belief to be true only by luck. For instance, I can guess truly that my mother will be at home at 2 pm, and form a corresponding belief. We are not prone to treat such lucky true beliefs as knowledge, because they are deprived of something important that prevents them from being knowledge. A *prima facie* reason to think that mere true guesses do not amount to knowledge is that the cognizer does not have a sufficient grasp of the truth of her belief, so she cannot be legitimately certain in what she beliefs. Whenever the truth of a belief coincides with the cognizer's hunch it cannot be a case of knowledge. Therefore, the lucky guess example shows that knowledge is more than a mere true guess. Some philosophers treat the surplus of knowledge over mere true belief in terms of some additional value that knowledge has over true belief. [20] Let us assume that they are right.

The appeal to explain the additional value of knowledge over true belief is known as the value problem, which I have already mentioned in chapter 1. Virtue epistemology and the KSA view is offered, and often promoted as an answer to the value problem. The proponents of the KSA view claim that knowledge is more valuable than mere true belief, because it is an achievement creditable to the agent. Call the thesis that knowledge is essentially an

[20] The intuition is traced back to Plato's *Meno*.

achievement the "achievement thesis" The achievement thesis will be in the core of our discussion in this chapter.

Duncan Pritchard (2010, pp. 37-41) calls the theory that takes achievement to be sufficient for knowledge "robust virtue epistemology". Pritchard offers two objections against this theory. He suggests first that we can have knowledge without achievement. Suppose, for instance, that Lili is in Budapest for the first time, and she asks someone on the street how to get to Kossuth square. By receiving correct directions, Lili acquires knowledge where Kossuth square is. But this knowledge is not primarily creditable to her.[21] It is primarily creditable to her informant. So, it is not her own achievement. Therefore, knowledge is not necessarily an achievement. Secondly, Pritchard claims that we can have achievement without knowledge. Such case is illustrated by the notorious fake-barn example. Suppose that Barney is walking in an area where unbeknownst to him there are predominantly fake barns looking exactly like real ones with few exceptions of real barns. Barney happens to look at a real barn and forms a belief that this is a barn. It is often claimed that his belief does not amount to knowledge, because he bumped into a real barn by sheer accident. Pritchard claims that Barney's belief is a cognitive achievement (it is primarily creditable to him), but it is not knowledge.

What goes wrong in these two examples? In the second one, Barney forms his true belief in a situation which is epistemically very risky. He could have easily formed the same belief, and be wrong. So, according to Pritchard, Barney has formed a lucky true belief similar to a lucky guess. As to the first example, Pritchard thinks that although what we get to know by testimony may not look like an achievement, nevertheless we gain knowledge by testimony because our beliefs so formed are usually safe – they could not easily be wrong. So,

[21] See Jennifer Lackey's example of Chicago visitor in Lackey (2009)

the moral of these two examples is that the achievement issue and the anti-luck issue are separate, and independent; and that the achievement thesis does not include in itself an anti-luck condition, and therefore it cannot be a satisfactory account of knowledge.

Pritchard suggests that even if robust virtue epistemology is correct in some respect (i.e. as a reply to the value problem) it has to deal with the anti-luck issue by separate means. His point is that the achievement thesis has to be complemented with an anti-luck condition in order to become a sufficient condition for knowledge.

In this chapter, I am going to argue against Pritchard's suggestion and in favor of the sufficiency of the achievement thesis. I claim that the crucial anti-luck condition is already integrated into the achievement thesis. Other kinds of luck are compatible with knowledge. I argue that the problem of luck has never been addressed in sufficiently fine-grained details, and the fact that there are malicious versus benign forms of luck has not been given sufficient attention. This resulted in a systematic misuse of the concept of epistemic luck. The majority of epistemologists tend to treat luck as being ubiquitously pejorative. They claim that luck rules out knowledge even in cases where luck is obviously benign and contributing to knowledge. My task in this chapter is to try to provide a sufficiently fine grained analysis of luck which will hopefully set a borderline between malicious and benign forms of luck. Then I will try to show that malicious forms of luck entail failures of success from ability.

The chapter consists of five sections. In the first section, I briefly remind the reader about the structure of the intelligent ability, namely that it can be analyzed into separate internalist and externalist grounds. In the second section, I discuss the principle of safety, and I distinguish between two notions of safety: normalized and counterfactual safety. I argue that normalized safety is the crucial anti-luck condition required for knowledge. I suggest that an intelligent skill has to be normalized safe in order to produce knowledge, and so success from

intelligent ability satisfies the normalized safety condition. Therefore, the anti-luck condition is included in the achievement condition. In section 3, I discuss safety of the internal ground of the ability of the skill. I claim that the necessary safety for knowledge amounts to having a safe internal basis as opposed to having that basis safely. I discuss the safety of internal basis as related to animal knowledge and to reflective knowledge. Finally, in section 4, I offer a detailed analysis of environmental luck – luck related to the externalist basis of the ability. I argue that construed in a certain way, positive luck is beneficial, and compatible with knowledge. In the last section, I briefly sketch my answer to Pritchard's objection that testimonial knowledge is not an achievement primarily creditable to the agent.

4.1.1 The structure of intelligent ability

Now, let me proceed to the details. The proponents of the achievement thesis think that knowledge is this:

S knows that *p iff S* believes that *p* truly on a basis *b* of her own intelligent ability that would produce truth in circumstances *c*.

Knowledge is supposed to be an achievement because on that theory it is thanks to ability or virtue that one manages to achieve knowledge. There are two main components on the right hand side of the biconditional: the internal basis of the ability and the external circumstances. The internal basis is constituted by a *skill, or virtue, or a combination between the two depending on the particular theory.* I take the internal basis of the ability to be an intelligent skill. The skill can be exercised in certain range of *circumstances* that constitute the external basis of the ability. The circumstances can be narrower or wider depending on the proficiency of one's skill. As I mentioned in the previous chapter, *S* can exercise her ability

only if she is in appropriate condition, and in a position to exercise certain propensities when necessary, such as attention, control, etc. For example, if one drives drunken, one might be incapable of focused attention, and hence she might fail to see a passenger on the road that she would normally see if not drunk. I am not going to address the issue of the condition in the forthcoming discussion.

The task until the end of this chapter will be to test how much luck is intuitively compatible with success from ability. I will explore this question in two parts. First, I will explore how much luck is compatible with the possession and exercise of the internal basis of the ability, the skill; and secondly, I will investigate how much whim on the side of externalist circumstances is compatible with them being appropriate for the exercise of the intelligent skill.

Let me make a preliminary remark. Taking both grounds of the ability separately and measuring them against epistemic luck can have possible bearing on our debate in chapter 2 where I argued that knowledge can be seen as prime and yet analyzable into internal and external components. Although I believe that ability can be analyzed into internal and external components I still hold that *exercise* of the ability is not merely decomposable into internal and external components. The next chapter offers a further explanation in which case we get a prime result and in which case we get a mere combination between the internal and external components.

4.1.2 Epistemic safety

So far we have seen that knowledge excludes certain types of luck. I roughly sketched the debate initiated by Duncan Pritchard, according to which the thesis that knowledge is an achievement does not adequately integrate the necessary anti-luck condition. In this section, I will discuss the adequate anti-luck condition that a belief has to satisfy in order to amount to

knowledge. One way to analyze luck – adopted by Sosa, Williamson, and Pritchard among others – is in terms of safety principle which I explain below.

The problem of luck is usually related to having adequate anti-accidentality conditions for knowledge. The internalist intuition about anti-accidentality concerns the right kind of internalist justification (having access to reasons for the belief), whereas the externalist construal of anti-accidentality links it to the stability of our epistemic powers to achieve truth.[22] The anti-accidentality condition that the KSA proponents adopt is typically externalist, as we have already mentioned in chapters 1 and 2. It naturally translates into a modal anti-luck condition, known as the safety principle.

Safety is a modal condition. Initially it was offered as an alternative to sensitivity principle, expressing the idea that one could not easily believe that *p* and be wrong, which formalizes as Bp>p. I will use a more recent, and refined version it (Dutant 2010) expressing the idea that a belief is non-luckily true if it is based on a method that cannot easily go wrong in close possible worlds. Let us use the safety principle as defined by Duncan Pritchard:

(SP) *S*'s belief is safe *iff* in most near-by possible worlds in which *S* continues to form her belief about the target proposition in the same way as in the actual world, the belief continues to be true.

It is crucial how we understand close possible worlds in order to asses exactly which possibilities of error are to be blocked by the safety principle. On a relevant notion of closeness, a possibility is close *iff* it resulted from slight variations of the relevant initial conditions. (Williamson 2000, pp. 123-124, Peacocke, 1999, p. 315). What we are going to do first is to present two possible interpretations of closeness.

[22] Here I mean the general epistemological distinction between internalism and externalism, not a distinction between narrow and broad competence.

Before I start with exploring the issue of closeness let me clarify two general things. First, I will be looking exclusively at the alethic modality: whether the way the belief is formed *de facto* produces true beliefs in relevantly defined close possible worlds. Secondly, although I agree that safety of a belief is necessary for knowledge, I will argue that Pritchard's safety condition is far too general and too strong. We do not need an extra strong such condition, because it disrespects the benign role of luck in certain cases. The form of the argument that I am going to pursue hereafter is that the safety principle needs to be more specific than the general safety offered by Pritchard. I will argue in section 4 that the relevant safety condition should block only malicious forms of luck, and it is narrower than suggested by Pritchard. Then, I will try to show that the specific safety principle, which I am going to defend, is integrated into the achievement thesis.

For the purpose of my argument, let me introduce a distinction between two notions of closeness, or two parameters that could be relevant to measuring closeness of possibilities of error: call them "counterfactual closeness" and "normalized closeness". The distinction has already been drawn in the literature (by Lewis, Hawthorne, and Dutant).

Here is a brief characterization of the distinction. Counterfactual closeness is a matter of what could in fact have happened, given the specific circumstances at hand. Counterfactually close possibilities are those around the particular space-time region in which the cognizer finds herself at time *t taken in a narrow focus irrespectively, or independently* of what could have happened in general given the default situation in which the cognizer has been settled so far. In other worlds, counterfactually close possibilities are those that could happen just now to the cognizer – we look at these possibilities in isolation from what typically can happen given the circumstances the cognizer's skills have been acquired and generally attuned to. The notion of counterfactual closeness is involved in the semantics for

counterfactuals[23] and is the one epistemologists have typically assumed. More precisely, epistemologists have typically assumed that in order to have knowledge, one need to rule out possibilities of error pertaining to the particular situation at hand.

In contrast, an alternative is normalized close to actuality iff it is a sufficiently similar variant of actuality that is at least as normal as actuality. Normalized closeness is construed as what could typically have happened in a class of normal alternatives to actuality, irrespectively of whether or not they could have happened in the circumstances at hand." (Dutant 2010, p. 2) The idea that the relevance of possibilities of error depends on their normality was defended by Alvin Goldman (1986, p. 107) and John Greco (2003, p.129–31).

Of course, the cognizer can just now be in a completely default situation, and hence there might be no counterfactually close situation for her which is not normalized close. In this case, the counterfactual and normalized closeness will overlap. In fact, this is what happens most of the time. However, the two parameters can go apart, and this fact, I will argue, seems crucial for the estimation of trickier cases that are standardly discussed in epistemology.

Now let us see how this distinction can be applied to the ability view of knowledge defended so far. We are interested in defining which possibilities of danger can prevent success from ability, and therefore knowledge. Close possibilities for exercise of a skill are those that would trigger an attempt to exercise it. The claim is that these possibilities can fall in two classes: counterfactually close ones and normalized close ones.

First, there are default possibilities for every cognizer. This is the set of typical circumstances in which one's intelligent skills are acquired and to which they are attuned.

[23] Counterfactual closeness is structurally expressed as a tree of branching time. Here is how Dutant describes this structure:" (A)t each time t we open parallel branches corresponding to slight variations of the conditions at t. A possibility p is close to w at t iff there is some sufficiently close prior time t' where a p branch opens. More generally,the later the last common node with a p branch is, the closer p is. The resulting closness measure appears to be the relevant one for evaluating counterfactuals. Call it counterfactual closeness." (Dutant 2010, p. 7)

These are possibilities that are generated by slight variations from her default cognitive situation. For instance, when I move from Bulgaria to Hungary, and I see a Hungarian forest the situation of seeing a forest is normalized close to my default situations of seeing forests. When I go to China, and I see a bamboo forest, this is a slightly more far-fetched normalized possibility for the exercise of my recognitional skills. A possibility of my seeing a talking forest is out of the scope of normalized possibilities for my recognitional ability.

The main characterization of normalized closeness is that it can be relatively far-fetched from one's current spatio-temporal situation. Counterfactually close possibilities, on the other hand, are never such that they could not actually happen to someone. This is the main difference between normalized close, and counterfactually close. Here are two examples of normalized close, but counterfactually distant. Take a BIV who has been well-trained to recognize physical objects on the basis of nerve stimulation. The possibility of her being embodied and presented with real physical objects is normalized close for her. But it is far-fetched from her actual situation. She might never been actually presented with real physical objects. Or take a more ordinary example: I may know how kiwi birds look like from Animal Planet movies, so the possibility of my seeing a kiwi bird is normalized close for me. But it can be far-fetched from my current situation. I may never actually travel to a country where I can encounter a kiwi bird. So, not every possibility that is normalized close is also counterfactually close. Compare that to an event that is normalized close but counterfactually distant. It is a normalized close possibility for continents on Earth to change their shapes. But it is not a counterfactually close possibility. It takes ages for the continents to change their shapes. A possibility that is far away in time is not counterfactually close.

Let me focus on normalized close possibilities in particular. I am going to use normalized closeness in a specific sense which is cognizer-bound. I want to refer to the default circumstances in which a skill has been acquired and to which it has been attuned as

the closest normalized possibility for the skill, and as determining other close normalized possibilities. I will stick to the term “normalized close”, instead of “normally close”, because our world as a whole can accommodate rarely encountered and thus abnormal regularities.

A skill can be well attuned to the frequencies displayed by the actual world as a whole. Most of our intelligent skills are meant to be like that. For instance, if Barney has such skill to recognize barns, his skill would enable him to reliably recognize barns in the world as a whole. Of course, the world has presumably a small number of abnormal regions containing barns that do not look like barns, or regions with predominantly fake barns. In the former sub-regions, Barney might fail to recognize that something is a barn, and in the latter sub-regions he might fail to recognize that something is not a barn. But still, if he manages to recognize a barn as being a barn in these regions, his belief would result from his skill to recognize barns, and not from sheer accident (unless he plays some special game of guessing). This is so, because these abnormal regions are parts of the actual world as a whole, and his skill is attuned to that world as a whole, so these abnormal regions are not excluded by his normalized circumstances.

Of course, a skill can be narrowly attuned only to an atypical sub-region of our actual world, and not to the actual world as a whole. Then only close possibilities to that particular sub-region would be normalized close for the skill. Such skill would be inadequate in relation to the actual world as a whole. In this case, the skill can have epistemic bearing only in that particular sub-region. Take for instance a professor who lives in a “black and white” sub-region of the actual world where there are only very good and very bad students. There, she acquired the skill to judge about her students’ papers by just quick skimming of the papers. The normalized circumstances for the exercise of her skill are narrowed down to her particular environment. They do not extend to the actual worlds as a whole where most of the students are neither too good nor too bad, but somewhere inbetween.

Nevertheless, some skills that are narrowly attuned to demands of a specific environment can be applied to the world as a whole. For instance, take an Eskimo who is exposed to predominantly seeing the white color of snow. Suppose that the Eskimo has developed a unique ability to recognize extremely many nuances of white in order to fulfill the everyday tasks of his environment. Although the skill of the Eskimo is narrowly attuned to his environment it does not interfere with his ability to recognize nuances of white elsewhere in the world. In this case the skill is attuned to the world as a whole, and has a potentially much wider application than just in his default environment.

We should resist the temptation to define normalized close possibilities as possible situations where were one to attempt to apply a skill, he could not easily go wrong. If someone tries to acquire a skill, but has not yet acquired it, the normalized close possibilities for the exercise of that skill would be the typical ones where one strives not to go wrong. Normalized closeness thus depends not just on one's default cognitive context, but also on her epistemic goals in this context. In other words, some possibilities might be normalized close for an intended skill, without the agent having that skill yet. It is very important to notice that normalized close possibilities are those for which the skill has been trained. The skill is trained or acquired for some cognitive purpose. The fact that there might be imaginary circumstances in which one's lousy strategies might work does not turn those lousy strategies into intelligent skill, because they are not actually trained for such circumstances.

Now, contrast the notion of normalized closeness to the notion of counterfactual closeness. Unlike normalized close possibilities, counterfactually close possibilities are ones around the concrete situation in which the cognizer finds herself. A good illustration of counterfactually close possibility is the possibility for Barney at time t to have looked at a fake barn instead of a real barn and to have formed a false belief that this is a barn. This is a possibility rather close to his actual situation at the moment t, and place p.

According to these two notions of closeness, we can define two notions of safety. A skill would be normalized safe iff it cannot easily go wrong in normalized close worlds. So, we can adopt the following principle of normalized safety of skills related to belief formation:

(NSP) *S*'s skill is normalized safe *iff* in most typically default possible cases in which *S* continues to form her belief about the target proposition in the same way as in her actual default world the belief continues to be true.

Analogously, we can obtain the notion of conterfactual safety from the notion of counterfactual closeness. A method of belief formation would be counterfactually safe *iff* it cannot easily go wrong in counterfactually close worlds. More precisely we can define this principle:

(CSP) *S*'s belief is counterfactual safe *iff* in most possible situations close to *S*'s current situation both spaciously and temporarily, in which *S* continues to form her belief about the target proposition in the same way as in her present situation the belief continues to be true.

We normally take the possession of a skill to be measured along normalized safety rather than counterfactual safety. The skill is expected to yield to success stably enough in a range of situations that are prototypically close to each other, centered on the default case. We never talk of a skill attuned to a particular spatio-temporal situation. It is always related to some default circumstances however narrow they might be. Therefore, counterfactual safety does not directly determine the possession of the skill. However, as I already mentioned, the scope of the circumstances in which a skill is normalized safe can be narrower or broader. There are skills that are tied to the actual world as a whole, i.e. safely yielding to success in the actual world as a whole, and skills that are attuned to, and reliable only in a sub-region of the actual world, but not reliable in the actual world as a whole. The first category of skills can produce knowledge not only in the typical circumstances in the actual world, but also in

the weird sub-regions of the world[24], where they are trapped in counterfactually unsafe circumstances; while the second category of skills cannot produce knowledge outside the close possibilities around the specific sub-region for which they are tuned, because they are not normalized safe in the world as a whole .

So, counterfactual unsafety does not have the power to deprive someone from her skill. If an archer manages to hit the target in still weather, but under the close threat of a wind blow, it is still going to count as success from her skill. It would be an instance of hitting the target on her credit. Some still weather conditions might be counterfactually unsafe, but most still weather conditions are typically not so counterfactually unsafe. Even if she fails to hit the target because of counterfactually close blow of wind that turns actual at the time of her shot, this would not take away her skill. This is so, because the normalized safety principle that I embrace entails just reliability, but not infallibility. Therefore, the relevant safety condition which determines the possession of ability is normalized safety. Now "success from ability" would be each case of reaching success via a skill (which is normalized safe) even if the cognizer's situation is counterfactually unsafe.

I suggest that we take the possession of intelligent skill to form true beliefs in a range of normalized circumstances as the norm of justification which provides the crucial anti-luck condition. Normalized circumstances, the ones for which one's skills have been trained have to be present, but they need not always be present in the counterfactual sense. Whenever one's true belief is reached via the skill, one's belief is justified and therefore it amounts to knowledge, since success from ability entails knowledge.

By adopting such a norm we get with it a theory of error. The norm is loose enough to allow for mistakes in some counterfactual circumstances. For instance, although being an excellent perceiver, Barney can easily take a barn facade to be a real barn, because he is

[24] Of course such skills cannot produce systematic knowledge in these sub-regions, but they are in a position to produce instances of knowledge. The idea is that these instances of knowledge would not be affected by

unaware of the specificity of his current situation. Of course this explains only a particular type of errors that we make.

In sum, normalized safety is a reasonably sufficient anti-luck condition which our true beliefs have to satisfy in order to amount to knowledge. This anti-luck condition is already included in the possession of intelligent skill. The skill should reliably produce truth in normalized close worlds which translates into - it should reliably produce truth in the typical circumstanced to which it is attuned. Because an intelligent skill is an acquired skill through practice which shares its structure with knowledge how (as I claimed in chapter 3) the true belief reached by such skill can be seen as an achievement. But since success from skill is an achievement, and the possession of skill entails the satisfaction of the crucial anti-luck condition, therefore the achievement thesis in my version incorporates the crucial anti-luck condition.

In what follows I will pursue the issue of luck in more detail. In the next section, I will focus on explaining in what particular sense the internal basis of the ability, i.e. the skill has to be safe. In the section afterwards, I deal with environmental luck. I will attempt to provide a detailed account of luck that the external basis of the ability can accommodate.

4.2 Safety of the internal basis of intelligent ability

4.2.1 Safety of the internal basis on animal level

Let us now start dealing with the question of how safe the basis of the ability should be in order to produce knowledge. As I said, the question has two sub-issues: 1. safety of the internal basis, or the method used in achieving knowledge, and 2. safety of the circumstances, i.e. the problem of environmental luck. I will start from the safety of the internal basis.

---

malicious epistemic luck that prevents knowledge.

When we talk of safety of the internal basis we are interested in the quality of the basis – we look at how it behaves in close possible worlds: whether it de facto produces true beliefs in close possible worlds. The main point in this section will be to make you appreciate a distinction between two kinds of safety concerning the internal basis. The distinction is between *how safe the basis is* and *how safely one has the basis*. I will argue that only safety of the basis is crucial for knowledge, but not how safely one possesses the basis.

Consider the following two examples:

Summy:

There are real cases of people who develop extraordinary abilities for a relatively short period of time during a stage of frontotemporal dementia. The phenomenon is known as 'the savant syndrome' (see Treffert D.A. 2009). Suppose that Summy, a subject of such dementia, develops an extraordinary ability to count big numbers. Call it "ability M". His ability can produce correct or successful results, and during its short existence it can actually participate in "apt" causal determinations.

Recovering John:

John lost his ability to recognize faces after a car crash. After an intense rehabilitation period he makes his first step in recognizing his wife Clara: he forms a belief B upon indeed recognizing his wife: "Here is Clara, my wife." Afterwards, for some seemingly long period of time he fails to recognize her, and then he recognizes her again.

There is no particular reason why Summy and Recovering John would lack knowledge. When Summy forms a belief that 7586.3540= 26854440 out of his extraordinary ability, the common intuition says that he knows the sum. [25] The same holds for recovering John. His belief B should amount to knowledge (since he actually recognizes his wife), although he does not have his recognitional ability safely.

The distinctive feature of these examples is that the resulting beliefs are based on safe internal grounds although the protagonists do not possess these internal grounds safely. In other words, they use grounds of such quality that cannot yield to false beliefs in close possible worlds in which the beliefs are based on these very same grounds. But on the other hand, in most actual situations in which the subjects attempt to apply these grounds they fail to have them, i.e. they do not have the grounds in a sufficiently stable manner. In particular, the basis of Summy's ability is safe: his beliefs cannot be wrong in close possible worlds based on that basis. Moreover, Summy's ability is at least minimally normalized safe, because Summy had had an ability to count numbers before he was overwhelmed by that new ability to count big numbers. It is just accelerated by intense brain processes resulting from inhibition of one zone of the brain that causes an accelerated activity in another. He actually lacked the ability to count that big numbers. The point is that does not have that particular ability safely, he might lose it tomorrow. The same holds for John. Suppose that he recognizes his wife when all neurons in his head fire properly. He could not fail to recognize her on that basis in normal conditions. But John's problem is that it does not happen to him stably enough when his wife is in front of him. So, he does not *have* his ability safely, although the basis that he uses in that particular belief formation is safe enough.

These two senses in which safety can be ascribed to internal basis have not been distinguished by reliabilists and virtue epistemologists. Reliabilists claim that a process is reliable not just if it cannot easily go wrong in close possible world, but also if it has an actual track-record of success (i.e. the process is reliably possessed by someone). The same holds even stronger for virtue epistemologists. An intellectual virtue is a feature (or ability) of one's cognitive character to reach truth in a range of circumstances. So, at the same time it has to possess the quality of yielding truth, and to be a part of one's character. These two senses can

[25] Greco (2000b, p. 182) shares this intuition, but Zagzebski (2000, p.208) tends to disagree that the 'idiot-savant' (who is permanently in Summy's condition) has knowledge. I find Zagzebski's intuition strange, since

go apart on my theory, and this is why my theory can account for why Summy and Recovering John have knowledge.

I am inclined to think of a properly firing neural network as a descriptive analog of safe internal basis. Just having a properly firing neural network can be seen as underlying the skill in a very narrow sense. I mention this, in order to show that my current intuition is not entirely speculative. However, I am not going to argue for this particular point here. The discussion in chapter 5 partly concerns the underpinnings of the skill.

It follows from these examples that a belief produced by a safe basis amounts to knowledge (at least to animal knowledge on Sosa's terminology) although *S* might not have the basis safely. This gives us a comprehensive explanation of the starting point of knowledge. I think the suggestion might also be an interesting alternative or complement to reliabilism, because now it is not reliability of the process or faculty in terms of track record that matter essentially, but reliability or even perhaps infallibility of its basis in a range of circumstances. All this holds for animal knowledge exclusively. Animal knowledge is crucially grounded in safety of a basis of an "animal" belief, but not on safety of having the basis.

4.2.2 Safety related to the internal basis on the reflective level

Remember Sosa's distinction between animal and reflective knowledge. Just for a reminder, Sosa claims that there are two levels of our epistemic ascent. On the first level our faculties produce knowledge spontaneously. He refers to that as "animal knowledge". On the second level, human beings are in a position to investigate the reliability of their first level beliefs and thus to obtain enlightening reflective knowledge. Reflective knowledge requires more than animal knowledge does. For reflective knowledge it seems that *S* must also *have* her animal abilities safely. This is so because reflective knowledge entails judgment about the

---

the savant has a sense of how things add together. He is not like a calculator in this respect.

reliability of our faculties on the basis of coherence. But again, it seems to me that one can have reflective knowledge when believing on a safe reflective basis (awareness of de facto coherence in a particular case), not necessarily on the basis of a stable coherent character. This holds especially for meta-knowledge related to particular instances of animal knowledge where the allover coherence of one's beliefs matter. For instance, one does not need a stable coherent character in order to know that one knows that a bear is chasing her if this is really the case. In this case, the subject implicitly takes into account that she is a reliable perceiver that she is not hallucinating etc. Of course, epistemic reflection requires more than just hitting upon a de facto coherent answer. If one uses randomly the right reflective procedure, this would not count as her having reflective knowledge.

So, I suggest that we use alethic modality once again for evaluation of reflective knowledge. Namely, we can just look at whether a reflective belief is based on a safe reflective basis, i.e. on *de facto* coherence with other beliefs, of course under the condition that *S* is aware of the *de facto* coherence. I don't want to reject that reflection has to do with reasons and that it essentially involves all or part of one's knowledge. If *S* is to reflectively know that p, *S* must be able to rule out possibilities that non-p, at least the relevant ones. In this case, one is to take into account evidence to the contrary of one's reflective beliefs. But we can just take as a main requirement a properly safe reflective basis on which the reflective belief is being formed.

Along these lines, Greco (2002, p. 299) claims that if the value of coherence is its reliability, then the difference between animal and reflective knowledge is in degree, rather than in kind. I prefer to say that the difference is in content rather than in principle. I suggest that we accept the internalist reflective access as providing a special kind of internalist content. This is how I see the way to accommodate epistemic internalism. But there is no special reason to evaluate reflective grasp in a principally different way than our grasp of

objects. Moreover, it is better to stick to one unitary criterion of knowledge. After all, reflection involves merely different kinds of intelligent ability.

To sum up, I presented the cases Summy and Recovering John to show that knowledge is possible if a belief is based on a safe basis without the subject having the basis safely. So, that the essential safety related to internal basis must be the safety of the basis, not the safety of having the basis. This holds both for animal and for reflective knowledge taken separately. Next, I turn to the safety of external circumstances, or the problem of epistemic luck most widely conceived.

4.3 External circumstances (Environmental safety)

The problem of epistemic luck has often been seen in the light of environmental luck. It has been thought that fake barn cases are typical cases where external circumstances prevent one from having knowledge. In this section, I argue against this intuition. My pro-luck argument partly hinges on the conceptual possibility of distinguishing between various forms of environmental luck. For instance, I take it that Barney can plausibly be claimed to be in a situation of "positive luck", as opposed to "ability preempting good luck". The distinction between positive luck and ability preempting good luck is crucial to my claim. It settles the boundary between benign and malicious forms of luck. It also divides Barney's cases from the Gettier cases by placing them on the opposite corners of the gap between knowledge and epistemic failure.

Epistemic luck can have various degrees, relevant in different ways to the evaluation of knowledge. This section is about defining more precisely what can we count as appropriate circumstances for the exercise of ability – appropriate to the extent that they do not preempt or block the exercise of a relevant ability. My arguments will be rather illustrative than conclusive. Sometimes I refer to appropriate circumstances as "enabling circumstances".

It is central to our task to appreciate some relevant degrees of environmental luck. I want to propose the following list starting from deficiency to richness of luck:

1. *Bad luck*. These are cases in which the cognizer is temporarily, or permanently disabled. Take for instance a piano player who lost a hand. Bad luck can be fatal in cases of permanent lack of ability, caused for instance by a disabling brain damage. Clearly in such cases one cannot have knowledge, so I am not going to involve them in the forthcoming discussion.

2. *Mitigated bad luck*. Mitigated bad luck is present in simulational environments which can be of very different kinds. Simulational environments are certainly beneficial for learning a skill, such as in the cases of astronaut's training, learning how to recognize Egyptian cats from pictures, or learning how to take care of a koala from a movie. But these environments prevent the exercise of ability, and hence reaching success from ability, simply because they are simulational.[26] This becomes clearer when we take BIV contexts. All BIV contexts fall in the category of mitigated bad luck. A BIV can obtain the ability to recognize Earthly objects, but might never be in a position to exercise that ability. This is why "mitigated bad luck" is in a sense bad luck, because knowledge is success from ability, and if one cannot exercise ability, one cannot have knowledge.

However, typically the contexts of mitigated bad luck are not cognitively misleading in relation to the simulated reality. They imitate reality that can as well be default for someone. So, in such environments one is in a position to obtain a skill that is "as if" normalized safe. It is only "as if" normalized because one's actual default context might not be at all like the simulated one; and also because some phenomenal properties of what it's like to have the ability at hand might be missing. Nevertheless, as if normalized safety can provide a good opportunity for obtaining quite sophisticated skill. In this respect, contexts of

[26] In the next chapter I distinguish between object-related ability, and appearance-related ability. I mean that simulation environments prevent that exercise of object-related ability.

mitigated bad luck can be epistemically better than contexts in which a skill is only counterfactually safe, but normalized unsafe, as the method of evaluating student papers by quickly skimming is good in some circumstances but fails in the default ones.

Let me connect this point to the previous chapter. I claimed there that knowledge-how is more fundamental than knowledge that. Contexts of mitigated bad luck are good illustrations of this point. Having the skill entails knowledge how to do something. This knowledge how, on its part, can underlie instances of knowledge that, because it is a skill that causes true beliefs in the presence of appropriate circumstances. In these contexts, one cannot have propositional knowledge, but one can have the more general knowledge how. The astronaut learns how to survive in the outer space. The BIV learns how to recognize external objects. And when it comes to the real success it is this knowledge how type that is responsible for causing it

3. *Absence of bad luck*. This is a situation in which appropriate circumstances for the exercise of the ability are present. Such situations are default, and they should be default if we are to have systematic knowledge. I call it "absence of bad luck", since in certain cases the ultimately full presence of appropriate circumstances is improbable. To illustrate what I mean, take for instance a case in which a pianist has to play a sonata on a piano which misses a key. Luckily, exactly that key is not needed for the performance of the sonata. In this case, performing the sonata is possible, and this very fact is a matter of absence of bad luck. Or take another example: full presence of enabling conditions for giving a public speech includes the presence of a microphone for the speaker. But sometimes a technical problem can occur even at important public speeches. So, the event can be interpreted as a case of absence of bad luck if no technical problem occurred. In our default cognitive contexts there is always some small risk of failure of knowledge in the described sense.

4. *Positive luck*. There are situations in which one hits the truth as a result of her skill to do so, although she could have easily failed to do so. The representative case here is the case of Barney. Another example is Sosa's example of the kaleidoscope perceiver that I outlined in chapter 2. This situation can be seen as a counterpart of the tolerable unsafety related to the internal basis: the case in which one does not possess the skill safely. Here, the unsafety pertains to the presence of enabling circumstances. When the cognizer does not have the basis safely, being it internal or external, her epistemic act is threatened with counterfactual unsafety: she could have easily failed to know. I will claim that counterfactual unsafety does not necessarily block the exercise of the skill. It does not block it when the agent is positively lucky.

5. *Ability preempting good luck*. This kind of luck is typical of Gettier cases. The Gettier agent, as suggested by Sosa (see chapter 2) has the ability, and she reaches truth or success but she does not reach it out of her ability. It is luck external to her ability but intervening in a way that blocks the exercise of the ability. Imagine an archer who skillfully shoots at a target, and he hits the bull's eye. However, between the shot and the hit there are two gusts of wind neutralizing each other. The archer arguably fails to hit the target as a result of his skill, but rather as a result of some miracle.

What I will do hereafter is to conduct a kind of thought experiment about our intuitive epistemological tolerance to epistemic luck. You can think of it as a test for our intuitions about the combinations between each member of the list and two other groups of parameters which I already mentioned above. I will examine how these varieties of luck apply to animal and reflective knowledge. I will also examine how they intercept with normalized and counterfactual safety. We have sixteen relevant combinations on a table including all the parameters. Each square including a possible combination is marked with a cross. The empty squares constitute impossible or highly irrelevant situations. Here is how the table looks like:

Kinds of knowledge	Normalized and Counterfactual Safety	Kinds of luck			
		Mitigated bad luck	Absence of bad luck	Positive luck	Ability preempting good luck
Animal knowledge	Normalized+ Counterfactual+	X	X		X
	Normalized+ Counterfactual-	X		X	X
	Normalized- Counterfactual+	X		X	X
	Normalized- Counterfactual-	X			
Reflective knowledge	Normalized+ Counterfactual+	X	X		X
	Normalized+ Counterfactual-	X		X	
	Normalized- Counterfactual+			X	
	Normalized- Counterfactual-				

I will discuss each combination against the assumption that the cognizer possesses a stable internal basis, unless one is in a train to obtain one. The first chunk of combinations applies to animal knowledge.

I start with a situation of mitigated bad luck that is both normalized and counterfactually safe. This is a context of simulation in which appearances of possible items are ordered in a sequence that seems to be both normalized and counterfactually safe. This is a sequence that represents a possible default reality for someone (normalized safe) and any particular presentation is coherent with the rest of the representations (counterfactually safe). In most such contexts real items are not present, but even if they are present, they are not visible to the subject. There is a veil of simulation, or representations of items. This situation can be illustrated by a simple example in which one learns how to recognize orchids from pictures. If the pictures correctly represent things in principle and in the context (i.e. there are no misleading tricks in the presentation), there is a great chance that one's true belief "This is

an orchid" when one sees a real orchid would amount to knowledge. However, in his current context, one does not have a chance to exercise one's ability; i.e. to actually recognize orchids, just because there are no actual orchids around. So, one can obtain a skill, and with it knowledge how, but one cannot exercise it. Such ersatz situations have epistemic bearing, if they correspond to the reality as it is, because from sufficiently representations/simulations that correspond to reality one can learn propositional truths or obtain cognitive skills which one can then apply to real authentic conditions. These are, indeed, the most usual contexts in which we learn something from books, pictures, words, movies, etc. For instance, one can learn where Madagascar is from a map. However, one does not have the opportunity to apply one's knowledge to real items, i.e. to find Madagascar by actually following the map.

Therefore, mitigated bad luck in this constellation can be a rich source of knowledge. One can obtain knowledge how, i.e. an intelligent skill, and also, one can obtain propositional knowledge, i.e. knowledge of facts. For instance, one can learn that Attila wanted to invade Rome from a movie. Note that classical BIV scenarios in which the BIV learns how from appearances fall in this category too. Such contexts allow for the subject to have knowledge, but also they can prohibit knowledge if the appearances do not present adequately a reality. This is why I put ersatz environments under the category of mitigated bad luck. In the next chapter, I introduce a distinction between object- related abilities and appearance- related abilities that is relevant for further consideration of such cases.

The next case presents our perfectly default cognitive situation: this is the case of absence of bad luck that is both normalized and counterfactually safe. This is a situation in which all enabling conditions are there. We have sufficient reasons to think that our default cognitive situation is often like that. This is a paradigmatic situation in which the cognizer is expected to achieve knowledge. However, I was careful not to call our default situation "full presence of enabling conditions", but rather "absence of bad luck". There is a danger of

systematic failure of knowledge lurking in completely full presence of enabling circumstances which I want to avoid. The danger may come from very friendly worlds, e.g. a world where *S* has a demon helper. In the same way as very hostile worlds, very friendly worlds are incompatible with knowledge. For instance, imagine a world, call it "TB world", where whatever one believes turns out true. Generic reliabilism entails that in such worlds the agent would have knowledge. Virtue reliabilists would most probably take it that in the TB world the agent cannot have virtues, and therefore she cannot have knowledge. I am on the side of the virtue reliabilists here. I share the intuition that knowledge is an achievement, which entails handling at least some danger of falsity. If in TB world *S* could have knowledge, knowledge would not have additional value over true belief. On the theory that I defended in chapter 3, in the TB world the agent cannot have knowledge, because intelligence cannot breed in such a world. TB world is epistemically over-hospitable.

This is my order-of-determination argument for the ultimate epistemic unsuitability of epistemically over-hospitable worlds: the only way for a world w to be overhospitable is that its states accommodate the hunches of the cognizer. This in a way inverts the canonical order of determination of cognitive-epistemic matters that runs from the world to the 'head'. The goodness of the world is dependent on enabling intelligent actions rather than enabling mere success. For such a purpose, it is important that the appropriate world is placed somewhere inbetween the friendliest one and the unfriendliest one. Our world, the world of absence of bad luck, is not typically overhospitable, and so it safes us the trouble of overhospitality.

The next square on the table is empty, because the presence of positive luck presupposes that the situation is counterfactually unsafe. Only when there is a close threat of failure, we can talk of success as being positively lucky. So, positive luck cannot be combined with counterfactual safety.

So, we move one square further to a situation that is both counterfactually and naturalized safe, but we have an instance of ability preempting good luck. Take the original Gettier example "Either John owns a Ford, or Brown is in Barcelona." Smith's method of forming this proposition is normalized safe, because he has enough evidence that John might own a Ford. So, he is entitled as far as normalized safety goes, i.e. by default, to form the belief that John owns a Ford. He could not easily form that belief without being right as far as he knows John's habit to always buy Fords. However, he is wrong about that particular Ford and instead Brown turns out to be in Barcelona. The last proposition makes the whole disjunction true. It is natural to construe this last proposition in terms of counterfactual unsafe. It would be counterfactually unsafe when easily Brown might not have been in Barcelona. Suppose we went there over the weekend. However, it might be the case that unbeknownst to Smith, Brown has got a new job in Barcelona, and he is staying permanently there. In this case, the proposition that brown is in Barcelona is counterfactually safe, and the whole disjunction too. Even in this case, we can have ability preempting good luck that undermines the epistemic status of Smith's proposition.

The next three examples concern a situation which is normalized safe, but counteractually unsafe.

First, we have a context represented by normalized safety and counterfactual unsafety combined with mitigated bad luck. This is a context in which normalized safety and counterfactual unsafety are only simulated, or just one of them is simulated. Matthias Steup (2004) provides an example of this sort. The example concerns a cognizer who correctly and reliably gets to the truth, but there is a nasty demon around who makes him feel unreliable, and hence uncertain. Now imagine that the way in which the nasty demon proceeds is to put a simulation cap on the head of the cognizer every time he forms a belief which presents the reality as if the object of the belief has disappeared which is not in fact the case. Does such a

cognizer have knowledge? On the one hand here she reaches the truth by means of her ability; on the other hand, her conviction in the truth is strongly shaken by the demon. Such cases constitute a challenge to the belief condition for knowledge, since it is hardly plausible that the cognizer would form a belief in such cases. So, I take it that this is not a case of knowledge.

The next combination of conditions concerns normalized safe, counterfactually unsafe and positive luck. It can be illustrated by the already mentioned example of Barney, a normal perceiver in the fake barn country, who happens to correctly form a belief "This is a barn". The same combination can also be illustrated by the kaleidoscope perceiver example drawn by Sosa which I outlined in chapter 2. As I already argued, such cases are cases of knowledge, because they represent success from ability. Barney has the skill to safely recognize barns in the normal world taken as a whole. The fact that he is right now in an abnormal region of that world does not take away his skill, neither has it plausibly prevented his correct recognition from being success from his skill, and therefore knowledge. Compare Barney case to the archer example. When the archer hits the target in still weather, but under counterfactual threat of a blow, we still count the hit as resulting from his skill, not as a matter of accident.

Here is how Sosa plausibly accommodates such examples in his theory. Sosa distinguishes between outright safety and basis-relative safety. The outright safety is the more general condition, namely that not easily would one believe *p* without *p* being true (B*p* > *p*). Most of our beliefs are not outright safe because there is always a threat that the purported perceiver may be dreaming. Hence, outright safety seems too strong, and Sosa argues for a more moderate principle, basis-relative safety, which applies to a belief when it has a basis that it would not easily have had unless true (B*p* has $\beta > p$). Basis relative-safety allows for having knowledge in counterfactually pretty unsafe situations. Sosa suggests that by claiming that (counterfactually) unsafe beliefs can still be apt:

> A performance can be apt though unsafe, in at least two ways. The performer's competence might be fragile, for one thing, and its required conditions might be endangered, for another thing. Despite these dangers, the performance might still be apt, however, though of course not safe. (Sosa 2007, pp. 92-93).

Finally, normalized safety and counterfactual unsafety combine with ability preempting good luck is a situation which typically creates the possibility for Gettier cases. Take again the original Gettier example where Smith believes that "Either John owns a Ford, or Brown is in Barcelona." As I explained above Smith's method of forming this proposition is normalized safe, because he has enough evidence that John might own a Ford provided by his default circumstances. However, he is wrong about that and instead Brown turns out to be in Barcelona. The last proposition makes the whole disjunction true, and this time it is taken to be counterfactually unsafe. It is counterfactually unsafe, because Brown is in Barcelona just for the weekend, and he might easily not have been in Barcelona.

Now we pass to the last combination in relation to our animal beliefs. This is a situation of lack of normalized safety and presence of just counterfactual safety. This situation can be combined with three kinds from our list.

First, it can be combined with mitigated bad luck. This would be a situation in which, for instance, *S* is presented with two pictures of airplanes, and she is asked to give an answer which one is Boeing 787, and which one is Pilatus PC-6. She has no idea which one is which, but she guesses rightly. Her belief is counterfactually safe, because in this particular situation, she could not form the same belief and be wrong. But she does not know, because she does not have the normalized safe skill to tell the right answer.

Secondly, normalized unsafety and counterfactual safety can co-exist with positive luck. This would be a case in which unbeknownst to *S* there are jokesters all around in her world. Suppose *S* forms a belief "The wall is white" in the rare moments during which no jokester is playing with light. Should we take these beliefs knowledge or not? I think we

cannot plausibly take them as cases of knowledge. This situation differs from Barney's case, or the regular kaleidoscope case in the default circumstances. While Barney is entitled to believe that this is a barn because in his world by default he could not easily form such a belief and be wrong, here *S* is not entitled to believe that the wall is white, because in her world is she to form such beliefs they would be wrong. She is not using a safe method on the first place. So here, the internal basis of *S*'s ability is undermined by the circumstances. Does this situation differ from the one where the individual learns how to recognize orchids from pictures? Again, the reader should wait until the next chapter where the latter cases will be discussed.

The final possibility concerning animal knowledge is the combination between normalized unsafety, counterfactual safety and ability-preempting good luck. As I said, ability preempting good luck is typical for Gettier situations. Although it is usually conceived that Gettier situations are situations that entail counterfactual unsafety this is not necessarily the case. It is possible to generate Gettier-type examples in counterfactually completely safe situations. But these would still be relatively abnormal, and thus normalized unsafe for most of us. Julien Dutant (2010) gives such an example which concerns a shy sheep case. His example concerns Sharon who, seeing a rock that resembles a sheep in the distance, comes to believe that there is a sheep in a field.There happens to be one, but it is hidden behind the rock. In fact, the sheep always hides behind this rock precisely because it has a sheep shape. Although the situation is counterfactually safe here, the cognizer has no knowledge in this case.

Now take the contrary situation. Suppose that Peter is very religious on Mondays, because his mother died on Monday. The rest of the time, he earns his living as a thief. Geri, who has a part-time job in the church on Mondays, meets him there every week. Upon seeing him she forms the belief "Oh, here is that religious fellow." Geri's belief is counterfactually

safe, and it is true at that very moment, but it is normalized unsafe. I think most would agree that such true beliefs do not amount to knowledge.

Let us turn now to reflective knowledge. Reflective luck can be evaluated in the same manner. Our first stop is a situation of mitigated bad luck that is normalized, and counterfactually safe, i.e. the simulation under question imitates the bit of corresponding reality correctly. This can be illustrated by two examples. For instance, during his training sessions, a future astronaut learns that she handles some tasks better than others. Another example is: going to an ophthalmologist to test her vision stereoscopically *S* learns she cannot see objects from certain distance. From one's performances displayed in simulational environment, one can learn how good one is in the fulfillment of certain cognitive tasks. So, one can indeed obtain reflective knowledge about one's own faculties in this combination of circumstances.

The next relevant situation corresponds to absence of bad luck. The situation here reiterates. The agent would normally have knowledge except in the TB world, where whatever one believes about one's faculties turns out true. For reasons that have already been mentioned, one cannot have knowledge in this case. It is not a case of success from ability.

Of course, complete environmental safety still does not prevent anyone from falling in a Gettier-type situation. Ability preempting good luck can affect our reflective beliefs too. Sven Bernecker (2006, p. 93) gives an example of that kind. Imagine that every time *S* is about to form a false perceptual belief, a demon helper arranges reality so as to make the belief come out true. The demon only cares about *S*'s perceptual beliefs and doesn't interfere when *S* is in the process of forming a veridical perceptual belief. Given the interventions of the helpful demon, all of *S*'s perceptual beliefs come out true. Quite obviously *S* cannot have reflective knowledge that her perceptual faculty is reliable, because she does not have a chance to objectively assess the reliability of her faculty. So, ability preempting good luck

rules knowledge out not just on animal, but on the reflective level as well. An important point that follows here is that while the demon-helper cases are counterexample to the general safety condition, they are not counterexamples to the multidimensional theory of safety defended here.

The next group of test cases involve normalized safety and counterfactual unsafety. The first situation on the table concerns mitigated bad luck when our reflective beliefs meet normalized safety and counterfactual unsafety. Remember that mitigated bad luck involves learning something from pictures, or simulations. So, this would be a situation in which, for instance a person, call him Rarney, is looking at a map of the fake barn country without knowing that it is a fake barn country. On the map, the few real barns are marked with a sign. Rarney wants to visit a barn and wonders what the sign means. He does not even think of the possibility that the unmarked barns are not real barns. Rarney thinks the sign marks the barn as being somewhat more special, so he decides to visit a barn marked with a sign. He is sure he knows he is going to visit a barn, and he is correct. However, his belief is intuitively not knowledge. Counterfactual unsafety is not appropriate for learning something, or gaining knowledge in situations of positive luck.

Next, we go straight to the positive luck, because we cannot talk of absence of bad luck in the presence counterfactual unsafety, where one's hitting the truth is made highly improbable by the circumstances. I am going to address, and contrast two situations here. The first one concerns the original example of Barney again. Imagine that together with forming R* "This is a barn", Barney also forms the reflective second-order belief R** that his R* amounts to knowledge. Since I took Barney's animal belief to be a case of positive luck, the first question is what kind of luck affects his reflective belief? Is it plausible that Barney's reflective belief falls under the case of positive luck too? Is his reflective belief "I know that this is a barn" a case of positive luck or is it a case of ability preempting good luck? Since it is

intuitively highly implausible that Barney can have reflective knowledge, I am prone to think his reflective belief would fall in the trap of a Gettier-type reflective situation affected by ability preempting good luck. Even if we assume that basis-relative safety on the reflective level is sufficient for reflective knowledge, and a reflectively safe belief does not need to be outright safe, Barney's reflective basis is simply unsafe. The basis of his belief is misleading, because many barns around appear as barns but are not.

Sosa (2007) is also inclined to deny that R** amount to knowledge, when he discusses the reflective knowledge of his Kaleidoscope perceiver example:

> That is a distinction worth deploying on the kaleidoscope example. The perceiver would there be said to have apt belief, and animal knowledge, that the seen surface is red. What he lacks, we may now add, is *reflective* knowledge, since this requires apt belief that he aptly believes the surface to be red (or at least it requires that he aptly take this for granted, or assume it or presuppose it, a qualification implicit in what follows)" (p. 32)

And he explains further why this is not the case:

> There, we fall short of reflective knowledge, however, because the jokester precludes the aptness of our implicit confidence that our perceptual belief is apt. His being in control makes it too easy for us to be confident in that default way, in normal conditions for the exercise of our perceptual competence, while still mistaken. So when, as it happens, we are right, not mistaken, this cannot be attributed to the exercise of our default competence as a success derived from it. (p. 36)

This passage tells us that there are two different reasons why R** cannot amount to reflective knowledge. First, something precludes the aptness of Barney's implicit confidence that his perceptual belief is apt. This violates the aptness condition. And second, it is too easy for Barney to be confident in the correctness of his first-level perceptual beliefs, while still mistaken. This violates safety of his reflective belief.

Let us see what reflective competence amounts to under Sosa's conditions. Reflective competence is the capacity for making a right judgement:

- about the normalcy of the circumstances (of perceiving)
- that one's first-level capacity is not defective, and

- that there is a relation of aptness between the perceiver's animal capacity and the truth of her animal belief, when the conditions of reflection are normal.

Barney is said not to have knowledge because he does not satisfy mainly the first condition: he cannot make a right judgment about the normalcy of the conditions (of perceiving). This is so, as I suggested, because the reflective basis of his belief is misleading.

We can still ask if there is a possible case in which one's reflective belief falls under positive luck, and thus amounts to knowledge. This would be a case similar to Barney's one, but of somebody who is in a position to make a right judgment about the normalcy of his particular epistemic conditions.

Consider the following thought experiment:

Quick switching thought experiment:

Suppose that since his birth, Little Prince (LP) has been envatted by an evil scientist. He is a capable, conscientious, and highly responsible cognizer, and he is excellent at what we would call 'recognizing roses from other flowers'. Suppose he forms a belief R: 'This is a rose'. The belief is false in the evil scientist's world and he does not have knowledge. Imagine that one day Little Prince is freed from the scientist's influence for a period of time, moved in an unconscious state into the actual world, where he wakes up in a situation subjectively indistinguishable from the initial bad situation. He awakens, sees a rose among other flowers, recognizes it and forms a token belief 'This is a rose', call it R#. Is his new belief knowledge?

Assuming that the scientist has systematic plans to simulate the real world (and presents LP with perfect simulation), there is no particular reason why LP could not obtain an ability to perceive in his simulated conditions. People can develop their abilities in simulated environments. Take an astronaut's on-land training. An astronaut, who has never been in the outer space before, can develop an ability to cope with the conditions there, once he gets there, although he has never actually been there, but only through a simulator. Similarly, LP

can develop an ability to recognize objects in our world by learning from his cognitive hallucination. So, given that LP is sufficiently adroit, his R* can count as success from ability and therefore animal knowledge. Also, LP can develop reflective competence in his bad world based on inductive evidence concerning:

- what are the conditions in which one is able to produce apt beliefs
- one's awareness about how one handles the task of first-order perceiving.
- whether one's first-level beliefs are accurate *because* adroit in normal conditions of reflection.

In a deontological sense, one's reflective competence consists in not disrespecting this evidence. Suppose that LP is perfect in fulfilling these tasks and hence obtains has reflective adroitness.

The challenging question is now does LP have reflective knowledge when transported to the new environment? Imagine that the professor had transported LP to the new environment under the condition that if he does not perform well there, he would be transported back to the vat for further training. Indeed, LP does well enough in his new environment, and *the professor is satisfied*. Now imagine that together with forming R#, LP also forms the reflective second-order belief R## that his R# amounts to knowledge. Can it possibly amount to knowledge too?

When transported to Earth, LP's reflective belief enjoys normalized safety, although his general situation is counterfactually unsafe, because the professor may decide to transport him back if he does not perform well, which would make his judgment about the normalcy of the conditions incorrect. His belief is normalized safe because it is formed in the typical environment to which his perceptual skills are attuned. His belief is counterfactually unsafe, because it is a close temporal possibility for him to be disembodied and transported back to his planet. However, his situation is not exactly the same as Barney's one. There is an

intuitive difference between the two cases that favors LP's reflective epistemic status over Barney's one. Unlike Barney, LP is in a position to make right judgment about normalcy of his conditions of perceiving. The professor is satisfied by his performance, and will not transport him back. Hence, the reflective basis that he currently has is not misleading in the way it would be if the professor were dissatisfied. Both his animal and his reflective faculties are attuned to our world, and now they can function well. LP is able to recognize that the normal conditions are there when they are there, but he is unable to recognize that they are not there, when they are not there. LP is finally in a situation of positive luck, and is able to do it. And since positive luck is compatible with knowledge, LP has even reflective knowledge. Again this is not the case with Barney who is not in a position to make the right judgment about the normalcy of the situation.

In sum, I try to reinforce the intuition that Barney has animal knowledge by offering a thought experiment concerning Little Prince whose brain was trained to achieve Earthly knowledge, and when transported to Earth for a brief period of time, he recognizes a rose, and thereby knows that this is a rose. In that particular thought experiment, more people are prone to think that LP has knowledge than in the Barney case. But this is unjustified because the two examples are exactly analogous. Both Barney and Little Prince form beliefs that are normalized safe and counterfactually unsafe. They both form their beliefs as resulting from their recognitional abilities in the presence of positive luck. They are not in situations of ability preempting good luck that is the typical context of the Gettier-type scenarios where the agent forms her belief from a false premise. Then, on the reflective level Barney falls in the category of reflective Gettier case. So, he does not have reflective knowledge. However, I tried to change that scenario for LP, and to show that it is possible to have even reflective knowledge in the presence of counterfactual uncertainty.

There is certainly a lot more to be said about each of these cases. My task, however, was to suggest a fine-grained theory of epistemic luck and to illustrate that we can coherently decide which forms of epistemic luck are compatible with knowledge by merely following our intuitive judgments about success from ability. Therefore, I think we can plausibly conclude two things. First, it is really normalized safety that matters for knowledge. And second, since this is established on the basis of our intuition about which cases count as success from ability (and therefore as cognitive achievement), it follows that the achievement thesis grounds an anti-accidentality condition, which I do belief is satisfactory.

4.4 The moral

I argued that the crucial safety is a status of the internal basis of our intelligent abilities. The internal epistemic basis (the method) must be such that it does not go wrong in typically normal situations. Normalized safety is supposed to incorporate the minimal subjective requirement for knowledge. An ability has to be normalized safe in order for the internal component to satisfy the requirement of subjective justification. Let me return and re-evaluate the clairvoyant problem at this point. Norman's clairvoyant capacity is not normalized safe, since it is not attuned to his normal circumstances. His new ability needs to adapt through adequate propensities invested in its exercise.

When it comes to the general issue of environmental luck, it seems that the commitment to the good actual world, which is neither too hospitable nor too hostile, is central to our epistemic enterprise. I attempted to show that on a relatively fine-grained analysis of environmental luck knowledge and epistemic credit (success from ability) turn out to be compatible with more varieties of luck than usually thought by epistemologists. For instance, it is compatible with some forms of environmental luck: *S* can know even when if *S* uses the same method that in normalized circumstances cannot go wrong, but in

counterfactualy close situation, it can be wrong. Note, however, that these situations are relatively rare encounters on a general scale of our actual world as a whole. Our analysis demonstrated that one is one's skills must be sufficiently attuned to one's default environment in order to be capable of producing knowledge. Presence of counterfactual safety is not always necessary. Knowledge seems compatible with positive luck in counterfactually unsafe situation. The completely full presence of enabling circumstances might not be compatible with knowledge as success from ability. They are not compatible when the world is such that it accommodates the hunches of the cognizer. Finally, the presence of mitigated bad luck puts the cognizer in a special situation, but it bears some contribution to her knowledge. The other species of "luck" from our list preclude knowledge. This holds both for animal and for reflective knowledge. I hope that my attempt to provide an analysis of epistemic luck gives us a reason to be more optimistic about the role of luck in epistemology, and more concrete about our anti-luck conditions.

In sum, although we need the presence of proper circumstances for exercising of ability, it is important to note that achievement is primarily creditable to the agent. It is not primarily creditable to the circumstances. If the theory is correct, some risk should be tolerable in achieving cognitive success, as far as the agent has done her job, and thereby has managed to reach the truth despite the close danger threatening her epistemic act. I hope that the argument has established that success from ability without knowledge is impossible contrary to Pritchard's first objection.

4.5 Is testimonial knowledge a case of knowledge without achievement?

Let me finally address Duncan Pritchard's objection that knowledge is possible without achievement because knowledge in the testimony case is not primarily creditable to the agent, but primarily creditable to her informant. First and foremost, I am committed to the

credit thesis in a slightly different way than virtue epistemologists are. I am not saying that knowledge must be creditable to the agent, but that it must be creditable just to one's intelligent skills which are not always shaped by agent's awareness. So a relevant objection against my version of the achievement thesis would be that what one learns by testimony is not creditable to his own intelligent skills. But such objection would be implausible for trivial reasons. Testimonial knowledge is obviously creditable to the intelligent skills of everyone who participates in the testimonial chain. At the same time, it may be creditable more to some individuals rather than to others. The point is that in case where knowledge is a group achievement, it is nothing but an achievement of each and everyone's intelligent skills.

The proponents of the achievement thesis standardly suggest that knowledge can be a group achievement. For instance, to use our example from the beginning of the chapter, Greco (2007) suggests that Lili gets credit for cooperative success; Sosa (2007, p. 95) suggests that she has partial credit which is sufficient for aptness, and therefore for knowledge. Wayne Riggs (2009) argues that we are not compelled to take every belief based on testimony as knowledge. For instance, if someone stops Lili and asks her where the Kossuth square is before she has gone there, she would hardly firmly assert that it is, say, four blocks away from here. She would probably not be so certain. So, she might not have gained knowledge by testimony if she feels uncertain about the trustworthiness of the testimony.

I think in each of these answers there is a part of the truth. Testimonial knowledge is a partial achievement of the agent, as Sosa claims, because the agent has to apply certain intelligent skills in understanding the shared information at least. This is enough for my achievement theory to go through.

Second, Greco is right that testimonial knowledge must be a cooperative success. Take the following negative example. In my childhood we played a game called "Russian telephone". The first one in a row whispers a word in the ear of the next participant, and the

word is spread around until the last participant. The last participant has to utter what she has heard out loud. The point of the game is to spread the word in such a way that every next participant has a poorer phonetic grasp of the word, and finally the word should end up funny. This is an illustration of how testimonial knowledge is not supposed to function.

Finally, the intelligent way to gain knowledge by testimony includes taking into account the trustworthiness of the informant. In this respect, we encounter cases of the kind described by Riggs. It is possible that I learn a location from someone on the street, whose credibility I cannot judge. In this case, I may not be certain where the place is until I reach it myself. My reaching the place as a result of the instruction then is not knowledge. I get to know where the place is after I reach it. However, most of us have the skill to judge the trustworthiness of the informants routinely, because the default situations are usually such that it is normalized safe to take testimonial judgments on face value unless there are indications on the contrary. But also people are generally sensitive to such indications. They avoid asking other people who look suspicious, and untrustworthy about locations unknown to them. This suggests that some sort of intelligent abilities even on the side of the cognizer are responsible, and indispensible for gaining testimonial knowledge.

4.6 Appendix: Table of combinations between parameters related to epistemic luck: my own evaluation

Kinds of knowledge	Normalized and Counterfactual Safety	Kinds of luck			
		Mitigated bad luck	Absence of bad luck	Positive luck	Ability preempting good luck
Animal knowledge	Normalized+ Counterfactual+	Yes. One learns how to recognize orchids from picture.	Yes, except in TB worlds		No. Shy sheep.
	Normalized+ Counterfactual-	No, Nasty demon.		Yes, Barney.	No, Typical Gettier cases
	Normalized- Counterfactual+	No, ability is too weak to produce knowledge.		No, Jokesters all around in my world.	No.
	Normalized- Counterfactual-	No.			
Reflective knowledge	Normalized+ Counterfactual+	Yes. Astraunaut training, vision checking	Yes, except in reflectively TB world	Yes, LP.	No, reflective Gettier cases (S.Bernecker)
	Normalized+ Counterfactual-	No, Rarney and the fake barn map. No, Nasty demon		Yes, Little Prince's global reflective knowledge	No, Barney's local reflective belief.
	Normalized- Counterfactual+			No	
	Normalized- Counterfactual-				

N.B. The "yes" answer indicates just possibility of knowledge in the given context. It does not necessitate it. The "no" answer indicates impossibility of knowledge.

CHAPTER 5: THE ABILITY VIEW OF PERCEPTUAL KNOWLEDGE

This chapter will explore the application of the view defended so far to the problem of perception. The main question we are going to attempt to answer is how our perceptual experience enables us to get in epistemic contact with external objects. The chapter has three parts. In the first part I try to defend the thesis that perceptual abilities play more essential role in shaping our perceptual experience than do perceptual virtues. In the second part, I discuss the role of perceptual abilities in the debate about perceptual knowledge. I offer brief reconstructions of the problems related to the traditional view of perceptual experience, metaphysical disjunctivism, and finally Alan Millar's arguments about the ultimate role that perceptual recognitional abilities can play in addressing these problems. In the last part, I discuss the theory of ability-based perceptual disjunctivism. I outline Alan Millar's view of perceptual recognitional abilities first. Then I present, and try to defend my own version of ability-based disjunctivism. I suggest that we notice an analogy between recognitional skill and knowledge how, and that this analogy has an explanatory potential that helps us to understand the role of recognitional abilities in bringing us in touch with the external object. I argue against Millar's understanding of ability as broad competence. I present a version of ability disjunctivism based on narrow competence.

The discussion requires a specific focus. The focus in a way cuts between epistemology and the philosophy of perception. The two fields involve different questions: philosophers of perception address questions such as "what is the nature of perceptual experiences?", "what is the object of such experience?", while epistemologists are interested primarily in whether we could achieve justified perceptual beliefs based upon our perceptual experiences or whether we can obtain perceptual knowledge thereby.

5.1 Abilities vs. virtues

Let me return to our debate in chapter 1. In Chapter 1, I offered three main arguments against virtues being necessary for knowledge, and I tried to show that all of them are compatible with taking intelligent skills to be necessary for knowledge instead. First, I claimed that reliable faculties are virtues only if they express virtuous epistemic character. But then, I argued, the assumption that knowledge results from virtues so construed is implausible. The reason why this is so is because knowledge is modular, i.e. some people are better in visual recognition than in mathematical calculations. If one is generally bad in mathematical calculations, one does not plausibly exercise a virtue when it comes to the domain of math. Nevertheless, even in this case one can have elementary mathematical knowledge: one can add, multiply, divide, and subtract. One can even manage to understand a more complicated theorem if one absolutely has to. However, since one is not reliable in mathematical thinking in general, and can easily go wrong one's mathematical knowledge cannot be plausibly said to be an expression of virtue understood as a reliable ability,. Therefore, possession of virtue is not a necessary condition for knowledge.

Also, I argued that the cognitive integration account of virtue is the most plausible one, because it explains how one's character can contribute to knowledge, and therefore comes closer to the idea that the agent as a whole contributes to knowledge. But this account still faces a problem. The problem is that the virtue of a cognitive system is not an achievement of the agent. If it were the agent's achievement, the agent must at least be aware of the motive for which the integration is being done. But she is often unaware of the motive that drives her cognitive integration. Sometimes her conscious epistemic motivation can even enter in conflict with her subconscious epistemic character that actually does the integration. Take a person who would rather not see scenes of cruelty of brutality. She cannot help seeing

them anyway if she faces them. In this case, seeing the scenes (an instance of veridical perception which presumably requires a subconscious cognitive integration) goes against the agent's conscious epistemic motivation. More precisely, it goes against the expression of weakness of her conscious epistemic character. This cannot happen with moral virtues though. A brave man would not regret entering into a challenging battle. Since, if he enters into a challenging battle and regrets it, he seems not to be brave. So, cognitive integration can certainly be a virtue of one's cognitive system as a whole, but it does not look to be a virtue of the agent. Another problem related to the cognitive integration theory of knowledge is that it is threatened by the conceptual possibility that an agent can know just one proposition *A*: "This is A". The agent has no other beliefs with which she can integrate *A*. If cognitive integration were necessary for knowledge, it would be unexplainable how the agent can know *A*.

My last objection was rather an illustration of the above arguments. It concerned the issue of perception which is particularly problematic for the virtue theorists. Perceptual knowledge is exceptional. Unlike its relatives, perception is automatic, passive, involuntary, and it is hard to say what kinds of refinements are particularly desirable for ordinary perception in the way that they are generally desirable for memory or reason. This makes its treatment in terms of virtues difficult, if not impossible. This is a central objection raised against virtue epistemology. For this reason, perhaps, there are indeed very few systematic attempts to account for perception as an intellectual virtue.

A different type of objection against the virtue theories was presented by the two examples that I gave in chapter 4, the examples of Summy and of Recovering John. In these two examples, the protagonists have at time t the relevant skills for reaching knowledge, although they do not have these skills reliably and therefore virtuously. Yet, the claim that Summy and Recovering John have knowledge when they produce true beliefs out of their

skills is intuitively very strong. Therefore, the examples show that virtues are not needed for knowledge.

5.1.1 Perception as intellectual virtue

Let us look in a little bit more detail at the claim that perception is not a virtue. As I said, perception is a rather convincing illustration why our faculties are not to be taken as virtues since it is automatic, involuntary and a subject to subconscious epistemic control that is not within agent's grasp. However, Stephen Neipher (2008) presents a systematic attempt to vindicate an understanding of perception in terms of intellectual virtues. He answers to a set of objections against perception being a virtue. The first objection which he discusses is "the unmotivated belief objection" (p.41). It says that perception is automatic, but if something is automatic, it is unmotivated. For illustration, my computer works automatically, but it does not do it out of a motive. We have the intuition that automata in general do not have their own motivations. On the other hands, virtuous actions are motivated actions. For instance, when one jumps to save a child from a tiger's hands, one is motivated by bravery and compassion. In the same way, when one opens Plato's dialogs, one could be motivated by intellectual curiosity. But when seeing a pear, one does not seem to be driven by any such motive for which the agent is praiseworthy. So, if perception is automatic and therefore unmotivated it cannot be a virtue. Therefore, perception is not a virtue.

Neipher replies to this objection by claiming that perceptual processing is driven by focused attention. Directing one's attention, he claims, is a function of motivations, because emotional processes play a role in directing one's attention. From this, he draws the conclusion that the kinds of motivations involved in directing one's attention whatever they are, are virtuous.

It seems to me that Neipher is too quick in drawing this conclusion. If motivations involved in directing one's attention were virtuous, they would be an expression of one's epistemic character. But I don't think that focused attention that contributes to knowledge is always an expression of agent's virtuous character. Take, for instance, the visual ambiguity of the Necker cube. The Necker Cube is a well known drawing of a cube presented by the Swiss crystallographer Louis Albert Necker, which looks like that:

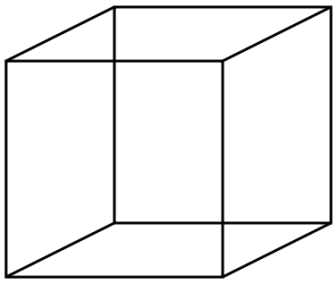

What is specific about this drawing is that the cube can be seen in two different ways depending on which side of the cube the observer takes as its front surface. The picture does not show which is in front, and which is behind and this makes it ambiguous. As a matter of fact, it is not always easy to switch from one to the other possible orientations even after one has seen the ambiguity. Doing the switch once entails focused attention, but it does not necessarily express a virtue – the person could be bad in doing the switch in general. So, he might not be in a position to apply his focused attention virtuously. Let me return to my main line that virtue is not necessary for knowledge. Notice that doing the switch once is enough to gain knowledge that the ambiguity is present. So the Necker cube is another example of knowledge which does not originate in agent's virtue; another example together with Summy and Recovering John of knowledge that results from a safe basis without the agent having the basis safely.

A different objection against perception as being a virtue is the involuntariness objection. It says that perception is involuntary, and therefore cannot be a virtue. Neipher's reply to this objection is that many virtues in general, such as the virtue of compassion are involuntary. A compassionate person cannot act otherwise but compassionately when she sees scenes of suffering, and in this sense the exercise of virtues is, and some would say, should be beyond one's control. So, the objection is refuted. Neipher's answer here is rather convincing. Yet, it is questionable whether the parallel between moral virtues and perception is fully sustainable. Compassion is an expression of excellence for which we praise the agent. But perception is usually not an expression of excellence for which we praise the agent. Of course, we can praise a musician for his ear, or a good painter for his eye, but we do not praise the ordinary perceiver for his eye or his ear. I want to stress just the point that we are not to praise *the agent* for gaining perceptual knowledge unless the agent is consciously engaged in directing one's attention, and thus gaining perceptual knowledge. In contrast, the musician and the painter are praised for the way they intervene or consciously mould what we hear or see. I am not denying that the agent's conscious effort can have an impact on what she notices and perceives. For instance, a loving husband can learn to notice the changes of his wife's hairstyle and hair color. However, the whole point I am trying to convey is that this is not absolutely necessary for gaining any knowledge by perception.

My line of thought here is perfectly compatible with the thesis, which I defended in the previous chapter, namely, that knowledge is an achievement of one's cognitive system. The eye and the ear, and the whole sub-personal setting of the agent can be entirely responsible for bringing up perceptual knowledge without the agent having any idea how this is being done, and also without guarding the process. The fact that we seem to see objects directly, and not though appearances or any other intermediate steps speaks in favor of the claim that the agent does not participate in these intermediate processes, but she just

consumes the final result. To put it simply, perceptual system can do the job of supplying us with perceptual knowledge. The agent's conscious effort might be helpful, but it is not absolutely necessary for gaining perceptual knowledge in all cases.

A third objection is that perception is unperfected natural disposition, but in contrast, virtue is a perfected natural disposition. We can call it "the perfected disposition objection". Neipher replies that perception is a disposition perfected through focused attention. Yet this is not enough to prove that it is a virtue creditable to the agent. Perceptual skill is acquired and refined through focused attention pretty automatically and sub-personally. The agent does not seem to have any idea as to what is going on during that refinement. So, the refinement itself is not to be credited to the agent. It is to be credited to her cognitive system.[27]

Therefore, we have no compelling reasons to treat perception as intellectual virtue. If we have no such reasons then the bad consequence for the virtue reliabilist in particular is that virtues seem to be parasitic on mere intelligent abilities (to perceive). Moreover, as I showed in chapter 3, we can have an equally beneficial normative epistemology based on intelligent abilities as that based on reliable virtues, so we do not need to involve virtues in our account of knowledge.

5.2 The problem of perception

The concept of ability has already been involved in the debate on perception (by Alan Millar). In a series of papers, Alan Millar (2008 a, 2008 b, in Pritchard, Millar, and Haddock 2010) claims that we gain access to the external object or state of affairs in virtue of exercise of perceptual abilities (discriminative and recognitional). Alan Millar presents a clear-cut

[27] Taking an already well-formed perceptual faculty, it is hard to say what kinds of consciously guided refinements are particularly desirable for ordinary perception in the same way as they are desirable for or faculty of reason. The only domain where such refinements would be relevant is art. But it is unclear whether perceptual refinements required for artists, or suggested in their artworks are really refinements in cognitive respect.

rationale for involving abilities in the debate. In what follows, I first introduce the problem of perception, related to the question of how perceptual experience put us in touch with external objects. This introductory part has a role of preparing the ground for a defense of the crucial explanatory role that perceptual abilities play in bridging between perceptual experience and external objects. Then I discuss Millar's view of abilities, and my own possible contribution to that view. My view differs from Millar's in that I take *possession* of ability to be compatible with absence of appropriate circumstances. I contrast the concept of possession, which I take to be narrow, with the concept of exercise of ability that is broadly construed.

The main question we are going to deal with until the end of the chapter is how perceptual experience puts us in touch with external objects. This question is tightly related to the question of how perceptual experience contributes to knowledge, which concerns both the issue of the nature of that experience, as we will see in a while, and the issue of justification of our perceptual beliefs on the basis of perceptual experience.

5.2.1 The source of the problem: indistinguishability between the good and the bad cases

Here is how the problems start. Imagine that Sarah perceives a cat. She has perceptual experience which makes her believe that there is a cat in front of her. Now imagine another case in which Sarah hallucinates a cat. Then Sarah has an experience such that there is no way for her to distinguish it from that experience that she had in the veridical case. She feels equally certain that there is a cat in front of her. Then, if these two states – good and bad – are completely indistinguishable for her, how can she be certain that there is a cat in front of her in the veridical case?

Both in epistemology and in philosophy of perception a lot has been written about the common nature of the states in which one perceives something, and states where one only

---

Without such refinements we are not exposed to making significant mistakes in our everyday knowledge in the

seems to do so, but which are in fact epistemically corrupted either by an evil demon, or by a mad scientist, or simply by hallucinations and illusions. On the basis of this reasoning, some seriously pessimistic conclusions about the nature of perceptual knowledge have been laid down, i.e. that perception like hallucinations and illusions provides no access to the external world, and therefore that we gain very little by perceiving. The trouble comes from subjective indistinguishability, which usually means – indistinuishability by introspection alone - of the good and the bad cases. Alex Byrne and Heather Logue (2008 p. 58) define subjective indistinguishability in the following way: a case* to be *subjectively indistinguishable from the good case iff*, in case* the subject is not in a position to know by "introspection" alone that he is not in the good case. Various theories of perception differ depending on how they construe the consequences of this fact of indistinguishablity.

5.2.2 Traditional interpretation of indistinguishability

The traditional view of perceptual experience starts from the fact of indistinguishability. Take an example in which Sarah is having a visual experience of a cat. According to the traditionalist, a case of seeing is a case in which: (a) Sarah has an appropriate experience, and there is a cat in front of her that causes that experience. According to the traditionalist, a case of hallucination is a case in which Sarah has an indistinguishable visual experience of a cat, and therefore the *same* experience, and the cat is not there. Holding this view of the sameness of the experience in the two cases, the traditionalist faces a kind of argument from illusion of the following kind. If the experience is the same in the two cases, it follows that whether a particular experience is perceptual or hallucinatory has nothing to do with its intrinsic character. Under the conditions of our example of seeing, the causal element cannot do the whole explanatory job of connecting

way we are exposed to making such mistakes in the domains of reason and memory.

Sarah's having the experience with the presence of the cat, because she can have that very experience without the causal connection. What remains mysterious is why she has the same experience without the causal connection. I claim briefly below that we cannot have a reductive account of seeing in terms of causal connection. The causal relation explains only partially how experience is connected to the object.

Having experience on the traditionalist kind is not sufficient in itself for seeing a cat. True, it is possible that the experience in both cases shares the most specific mental kind, and yet be epistemologically different, because in the veridical case the relevant object causes the experience which does not happen in the non-veridical case.

The first reason why the causal relation from the object to experience is not sufficient is rather simple. The object causes only raw stimuli on the retina. An infant without the suitable equipment of concepts and recognitional skills can only vaguely discriminate these stimuli; she is not in a position to recognize objects. This is a case in which despite the fact that her perceptual experience is caused by the physical objects, it does not amount to experience of those very objects. In order to be able to have a veridical experience one needs to build up a perceptual skill. So, it is not true that a simple causal relation from object to experience can account for that experience being an experience of these very objects. We also need perceptual skills to shape the experience in a way that it be a perceptual experience that adequately grasps the objects.

Second, the view that the traditionalist experience when caused by the external object amounts to veridical experience faces a problem that all causal accounts face - the problem of deviant causal chains. Imagine that I have a hologram in my room which presents an indistinguishable image of an object only if that very object is in the room. In fact I use it to detect whether my rabbit is in the room. Imagine that my rabbit is hidden behind the wardrobe in the room when you enter the room. You look at the hologram and you have an experience

as if the rabbit is in front of you. In this case the rabbit causes the hologram to display its indistinguishable appearance, and it is the hologram that causes your experience. But by transitivity of causal relation, it is the rabbit that causes your experience. However, you do not see the rabbit. Therefore, a reductive account of seeing in terms of causal connection is not possible. We need something more, an additional component which would differentiate the experience in the good and in the bad cases.

5.2.3 Metaphysical disjunctivism

Shall we opt for a different theory of experience? Disjunctivism about perceptual experience (also known as metaphysical disjunctivism)[28] is a different theory about the nature of veridical experience which aims at solving the theoretical problems related to the traditionalist account. Disjunctivism is the view that the most specific mental type of experience in the good and in the bad cases is not the same, or in other words, that at least some bad cases are mentally radically unalike the good case.

Why disjunctivism? The idea that an exclusive disjunction divides the two cases was brought forward by Hinton, whereby the name "disjunctivism", although Hinton does not use the label 'disjunctivism' for his own theory. The coinage is due to Howard Robinson (1985). The disjunction was expressed in different ways by different authors: for instance, that there is no "kind of experience common and peculiar" to the good and the bad case or that there is no "common element" (Hinton 1973, p. 62); no "highest common factor" between the good and the bad state. (McDowell 1982/1998: 386); no ''single sort of state of affairs'' obtaining in good and bad cases (Snowdon1980–1: 186); or no ''distinctive mental event or state common to these various disjoint situations''(Martin 2004: 37). Disjunctivists, of course, assume that the two states are indistinguishable but they take this fact not to play a

[28] For a presentation, and classification of various kinds of disjunctivism, see Byrne and Logue (2008),

constitutive role in determining their mental natures. Metaphysical disjunctivists vary as to how different the experience is in the two cases is: where more radical ones argue for a phenomenological difference, while the more moderate views involve difference in some other mental respect.

How does this difference bear principally on answering the question of our concern? Disjunctivists understand perceptual state as a state in which perceptual experience encompasses or internalizes the external state of affairs. For instance, Sarah's seeing a cat encompasses the presence of a cat. Hence, the answer to our question would be: experience puts us in touch with the external object by encompassing it. What happens during hallucination, in contrast, according to the disjunctivists, is that the state of perceptual experience encompasses just an appearance. In this essential respect, experiences that we have in the good and in the bad cases are very different. On the relational conception of experience (J. Campbell), perceiving *p* entails having an essentially relational experience, such that it would not be the same experience unless *p* is present. For instance, when I see a book on the table in front of me, my experience is such that it would not be that very experience unless the book were there on the table.

The relational conception of experience provides a framework in which we can see experience as putting us in direct touch with the object in the good case. Thus it overcomes the problem of deviant causal chains that the traditionalist has to face. When you see just the hologram in my room, you do not have a veridical experience. You have to actually experience the object. However, it still faces the first problem that we raised against the traditionalist view. This was the problem from an infant's experience. The object causes only raw stimuli on the retina. An infant experiences the object as raw stimuli although its experience puts her in touch with the object. But it is not a veridical experience of the object

without the suitable equipment of concepts and recognitional skills can only vaguely discriminate these stimuli; she is not in a position to recognize objects.

Paul Snowdon (1992) suggests that in order to explain the nature of veridical experience it is of primary importance to explain how demonstrative thought is possible. He takes it that the veridical and the non-veridical cases provide different truth conditions for the corresponding "looks" judgments. In the case of perception, Snowdon thinks, the truth-conditions are constituted by some feature of the relation to the object. In the case of hallucination they are constituted by the non-object involving inner appearance. In other words, Sarah's judgment "This is a cat" is made true by her relation to the real cat. In the case of hallucination, her judgment is supported by the mere appearance of the cat. Demonstrative judgment is determined not only by the relation between the experience, and the object in the good case, as the causal theory have it, but also by the nature of that same experience, i.e. the fact that it stretches out and encompasses the object. We can say then that an infant's experience cannot satisfy the truth conditions of demonstrative thought that veridical experience usually satisfy, so this is why it does not put the baby in touch with the object.

Now, the thought that the experience extends to, and involves the object can be found strange by the opponents of disjunctivism. They might argue that the step from the traditionalist view to the claim that experience has different mental nature in the good case is somewhat ad hoc. In particular, they can reason as follows. The causal theory raises some problems, say, of deviant causal chains, but if the disjunctivist assumes that the experience extends to and embraces the object, then deviant causal chains are blocked. So, the disjunctivist move is offered just as a solution to the problems of the traditionalist view, and indeed an implausible one. Experience is internal and private. It does not get out of the subject to encompass the object.

5.2.4 The role of perceptual abilities

I think, however, that the disjunctivist claim has a *prima facie* plausibility. Yet, it seems to me that one reason for that is because there is a convincing story to be told in its favor. The story can be provided by reference to perceptual-recognitional abilities. The idea would be that experience links us to the external object by means of exercise of recognitional abilities. Alan Millar introduced that idea. He claims that bringing perceptual abilities into the debate on perception clarifies the metaphysical statement that one's experience reaches out and encompasses the object. Take the example of Sarah again. In the good case Sarah exercises her recognitional ability of seeing a cat, and thereby knows that there is a cat in front of her. In the bad case she does not exercise any such ability, and therefore fails to know that there is a cat.

I want to offer a further step: perceptual abilities play a role of perceptual justifier. My belief that *p* is justified by the fact that I perceive *p* by exercising a recognitional ability, and therefore, my perceiving *p* can serve as a reason for justification of other beliefs or actions. Thus, the epistemic standing of perceptual beliefs is dependent on the exercise of perceptual recognitional abilities. The contribution of recognitional abilities to veridical experience explains also why our infant who perceives just row stimuli does not have a full fledged veridical experience. She simply has not yet developed recognitional skills.

The whole point of involving abilities in the debate on perception is to show that they do the main job of linking experience to the external object. Let us look at a possible explanation of how exactly they do that. A fully contentful explanation can be provided by psychology, although of course a normative generalization of it is possible. Assume that the main mode of mental presence in perception is indeed focused attention, as Neipher suggests. Indeed, human eyes move to allow inspection of details of the world, visual functioning depends critically on the individual's ability to fixate and to make saccades and pursuit

movements. An infant has this mode as an inborn function, yet she is not born with discriminative and recognitional skills, she cannot discriminate and recognize objects immediately. The infant is simply born with discriminative and recognitional capacities, which are there to be transformed to skills by the relevant propensities, in our case - the mode of focused attention.[29] I take it that having the capacity is having a natural and unrefined disposition, whereas having the skill is already having the refined intelligent ability. Also, I take it that the relevant mode, or propensity, e.g. focused attention partly underlies the capacity. [30]

So, we can assume that, for example, in its initial stage the visual capacity allows us to see just raw stimuli, but after it is being trained trough focused attention it grows into a recognitional skill which enables us to mentally reach the objects. The specificity of recognitional skill is that when we exercise it we see object directly, and not through the veil of stimuli. The idea is that the senses function from birth, but they do not reveal a world of objects initially. They reveal only a vaguely ordered picture of meaningless stimuli. Grown up perceptual experience is something different, and more than that. As *Gestalt* psychology showed, perception involves seeing patterns, and interpreting patterns that are more than the mere stimuli. Thus eventually, seeing via recognitional abilities is something rather different from what one actually and literally receives on one's visual field, i.e. from the mere sensing. Via the exercise of a well-developed recognitional skill, perceptual experience presents us with a world of objects, and not with a world of chaotic stimuli. Likewise, with other perceptual modes different from seeing. In this spirit, Millar (in Pritchard, Millar, Haddock 2010, p. 143) argues that once we have the recognitional skills, their exercise does not entail inference of objects and patterns from the stimuli, but directly seeing them. In this way, he claims, a theory of recognition accounts for the phenomenology of perception.

[29] I assume that the infant has a healthy brain.

Mohan Mathen (2005) develops a philosophical theory of perception along these lines, which is a nice supplement to the above claim. Mathen distinguishes between perceptual experience and perceptual exploration. While experience is passive and about sensations, exploration is active and about external objects. Perceptual exploration, according to Mathen, is searching for the cause of one's experience which consists of mind-independent properties, and objects. He categorizes perceptual exploration as 'a low-tech skill', like knowledge how, and unlike scientific methods. The full process of exploration gives us empirical certainty about things in the world. Indeed, Mathen shares the view that exploration shapes the experience in a way that makes it an experience about external objects, and not about the perceiver's sensations.

What is happening during the refinement and transformation from mere perceptual capacities to perceptual skills would provide an expected explanation of how our perceptual abilities contribute to knowledge, and why they have to be credited for that. But this is not a story to be told exclusively by a philosopher. As I said, for a fully contentful story we should look at the literature on perceptual development, or to engage in a discussion with cognitive scientists.[31] However, for our purposes it seems safe enough to assume that on the brain level the formation of skill entails the formation of a robust neural network(s).

The point of our direct interest here is that the acquisition and the exercise of perceptual skill happen without the agent's awareness, and so it cannot be properly called a virtue of the agent. Instead, as I claimed in chapter 3, the refined ability resembles closely the structure of knowledge how as presented by Ryle, and it is refined through propensities. Appreciating the role of guiding attention (the special kind of perceptual propensity) in both the formation and the exercise of recognitional skills shows that perception can be seen as a

[30] If we go naturalized, focused attention can be seen as the right propensity selected through evolution. The rest of the capacity would include the brain and personal DNA dispositions.

species of knowledge how. This is, I suggest, a more robust alternative to Neipher's claim that focused attention reveals that perception is an intellectual virtue. Notice that we say "I know how to recognize x, or I know how to discriminate x", despite the fact that it is not natural to say "I know how to see" in most contexts. If perceiving is clearly based on discriminative and recognitional skills, though, this suggests that perception must have something in common with the rest of our intelligent abilities that are structurally alike knowledge how. Appreciating this kinship can be helpful in understanding the job of recognitional skills in linking our experience to external objects in analogous ways to the jobs that other practical skills do in "hitting their proper targets" as shown in chapter 3.

5.3. Ability-based perceptual disjunctivism

5.3.1 Millar's theory of perceptual abilities

So far I presented one possible viewpoint on the problem of perception. We saw that the traditionalist view might face an argument from illusion if it relies exclusively on the causal relation between the perceptual object and the experience to explain how veridical experience puts us in touch with the object. Metaphysical disjunctivism overcomes the difficulties of the traditional view, but some traditionalists might be dissatisfied with the disjunctivist story. In particular, they might find it unbelievable and ad hoc. I suggested that one way to back the disjunctivist answer is to outline the role of perceptual recognitional abilities. In what follows, I will first introduce Millar's view of perceptual abilities, and I will argue against the need of broad competence in the ability theory of perception. In the next section, I will defend a version of disjunctivism based on narrow competence.

[31] For instance, see Alan Slater (Ed.) 1998

Millar sees perceptual abilities as falling in two groups: discriminative and recognitional abilities. He claims that they are both akin to practical skills, and as I suggested we can actually treat them as a species of knowledge how. Discrimination, he thinks, is necessary for recognition, but he prefers to stress on recognitional abilities, because recognitional abilities account for phenomenological immediacy of the perceptual object. Millar defines recognition as the ability to tell by sight that something is F, which is opposed to the notion of perceiving as inferring from superficial features, or from how things appear. By means of recognition, we grasp an object in an immediate way no matter how its particular features are being discriminated in *S*'s visual field. In recognizing, the perceiver sees immediately a *Gestalt*. Thus, the theory of recognition accounts for phenomenological immediacy of perception.

An argument that Millar uses in favor of his theory of direct recognition, and against the view that recognition is an inference from features is an argument from underdetermination of perceptual descriptions. Imagine that I am waiting in front of a professor's office. You approach me, and ask me whether there is a student in the office. I say that yes, but that I don't know his name. You ask me how he looks like. I say that he is an average tall dark-hair guy with glasses. There is a pretty good chance that you won't recognize who is exactly in the office on the basis of my description, although he might happen be a close friend of yours. Millar refers to this phenomenon as "underdetermination of perceptual descriptions". He uses it to show that we are generally not good in figuring out the objects/subjects from sets of typical features, and thus to undermine the theory that perception is an inference from superficial features. The point might be that if we cannot verbally describe the content of perceptual experiences, then presumably we cannot so quickly conceptualize premises for the purported inferences.

As a matter of a different fact, cognitive scientists have produced plenty of computer programs for recognition based on the principle of inference. All of them have proven to be far too slow in comparison to human's average speed of recognition (see Norman 1993 and the whole issue of this journal). Having in mind the fact that computers have much more powerful capacity to calculate than humans do, makes it very improbable that the underlying mechanism of human recognition is actually inference from features. This fact from cognitive science is relevant to our philosophical theory that recognitional skill enables us to see objects not by inference from features but rather in a direct fashion.

By accounting for the phenomenological immediacy of perception, recognitional abilities supply metaphysical disjunctivism with additional backing. Recognition links one's mental state to the external objects directly, and not via appearances (i.e. it is a tool for taking the fact directly in one's consciousness, and not taking an appearance in one's consciousness). For instance, a coin may look elliptical to me from most angles of sight, but I immediately recognize it as a coin, and therefore as circular. [32] This, it seem to me is a plausible story to be told as to how one's mental state encompasses an object or brings a fact within one's consciousness.

5.3.2 Narrow vs. broad competence

There has been a recent debate about narrow vs. broad competence.[33] The theory of narrow competence claims that ability is analyzable into internal and external components (grounds), and the possession of the ability is identified with the presence of internal ground. The theory of broad competence, on the other hand, claims that the actual presence of proper environment is necessary for the possession of the ability. For instance, it is claimed that one is able to recognize water only if there is stuff with chemical structure H2O is generally

[32] Jesse Prinz (2006) defends the claim that we can perceive things that we do not sense.

present in her environment. She lacks the ability to recognize water if a stuff with structure of XYZ exactly indistinguishable from water is predominant in her environment instead of H_2O. The theory of perceptual abilities that Millar defends is a theory based on broad competence. (in Pritchard, Millar, and Haddock 2010, pp. 167-174). He thinks that the possession of recognitional ability is essentially determined by the presence of appropriate environment for its exercise. Such environments, he claims, are environments where appearances of things are representative of things themselves. Millar and Williamson, who are both proponents of the broad competence theory, think that ability is not analyzable into internal and external grounds. However, a theory of broad competence is in principle compatible with an analysis of ability into internalist and externalist components. In this version, the theory would have it that the possession of internal ground is necessary, but not sufficient for the possession of ability. It won't be entirely constitutive of the ability.

The internalist (narrowist) opponent view is the view that abilities are constituted by separate internalist bases, and abilities can survive unfriendly environments. The internalist would claim that a person who has the ability to recognize water would essentially keep it in the XYZ environment, but would not be in a position to exercise it in that case.

Now if we want to treat recognition as the tool for extending one's mental state to the object and grasping it, the act of recognizing is a broad state. However, it is a separate question whether recognitional competence is broad. A defense of a theory of knowledge based on ability requires us to address this issue.

A quite representative case of the debate is the case of Barney again. On the broad theory, Barney loses his ability to recognize barns in the fake-barn country; on the narrow theory, Barney retains his ability to recognize barns even when he is in the fake barn country.

[33] For instance, see Williamson (2004) vs. Sosa (1993).

The proponent of broad competence would claim that Barney cannot have knowledge as to where there is a barn in front of him; whereas the proponent of narrow competence would claim that Barney can still have knowledge when he looks at one of the few real barns and forms the corresponding belief.

One important factor in the debate of narrow vs. broad competence is how one is prone to use the word "ability" or "competence". I cannot resist the following consideration here. In my native Bulgarian language the notion of ability cannot be used in a broad sense. One cannot say in Bulgarian "I am *able* to swim" meaning that there is a swimming pool around. Unlike in English, in Bulgarian this phrase means strictly that one has the skill to swim, and not that one has the opportunity to swim. In Bulgarian "to have an ability" means strictly to have a capacity or a skill. If I want to say in Bulgarian that there is a swimming pool around, I should say "I have the opportunity to swim", but not "I have the ability to swim". This is why I am inclined to think that ability is a narrow concept.

So, I am naturally inclined to defend a theory of narrow competence. Hereby, I would like to suggest a separate consideration in favor of narrow competence by raising an objection to the theory of broad competence. A problem related to broad competences may come from cases of learning in simulational environments. Millar (in Pritchard, Millar, and Haddock 2010 pp. in 162-167), for instance, claims that one can have recognitional ability only in environments where the appearances of things are distinctive of things that are actually there. There are environments, however, where the appearances of items are distinctive of the items, although the items themselves are not present. Here are two examples:

Astronaut's training:

An astronaut is being trained to survive in the outer space. For the aim of the training, he is put in a simulation environment where the simulated features are distinctive of outer space. However, he is not in outer space. Does he learn, and therefore obtain an ability to survive in outer space from the simulational environment? It is pretty clear that he does. Otherwise, it would be a miracle how he is able to survive after he has been sent into the outer space.

BIV in the hands of a friendly scientist:

Imagine a scenario in which all humans capable of knowledge are being exterminated, except for a scientist Prof. Kind, who survives by luck and meanwhile manages to envat the brain of a baby Little Prince before his brain death. With the good care of Prof. Kind, both survivors safely get to a peaceful planet where Prof. Kind dedicates the rest of his life to train the Little Prince's brain. He teaches him to recognize objects in our world by feeding him with coherent hallucinations. The brain is brought to think that he is an embodied little boy, who moves and grows freely, who enjoys the contemplation of one rose, etc. Prof. Kind hopes that one day the brain will be safely transported back to our world, embodied, and the new Little Prince will be capable of (or competent in) producing human knowledge based on his careful and conscientious brain-training. All appearances that the Little Prince is fed with are, and are carefully meant to be distinctive of objects on our planet. However, no such objects are present in LP's present environment. Does Little Price learn the ability to recognize Earthly objects in his environment where no appearances is representative of any object? If everything goes fine with the professor's experiment, LP must have the competence to recognize Earthly objects.

What the theory of broad competence can hardly explain is the epistemological significance of such environments. Since, on that theory, the possession of the ability is inseparable from its proper environment, the question remains how come environments that

are entirely inappropriate for the exercise of ability can enable one to obtain the ability? My argument is perhaps not compelling against the theory of broad competence, but it strongly suggest that ability can be analyzed into separate internal and external components, and that the external components can be graded in their epistemological bearing towards the internal component as shown both by the above examples and in the second part of the previous chapter. I mean that obviously some environments can contribute to the perfection of the skill, others to its exercise. Generally, the theory of broad competence allows for taking the skill, or some other internal component as being necessary but not sufficient for having the ability, as I already mentioned. The broad theorist can argue that the skill (or some other internal component) but not the ability as such, can be acquired in simulational environment. Alan Millar, for instance denies that abilities can be analyzed (in Pritchard, Millar, and Haddock 2010, 174-183), but at some places, Millar (2008b) takes it that abilities have some sort of internal grounds, i.e. reliable modes for belief-formation tied to the appropriate environment. He claims that such modes are acquired in simulational environments. However, it seems to me that there is a remaining controversy in his theory. How these modes differ from narrow competence, and why they can be acquired in totally inappropriate environments such as the BIV environment, and the astronaut's simulational environment, remains unclear to me.

5.4 Ability disjunctivism based on narrow competence

In the previous section, I have argued that in order to play the purported explanatory role complementing metaphysical disjunctivism, a theory of perceptual ability needs to be internalist. It has to ascribe to ability a constitutive internal element responsible for putting perceptual experience in touch with external object. I formulated such a theory in chapters 3 and 4. Now, if I want to argue that recognitional ability is narrow, I have to answer a further important question. Is there a way to make the internalist intuition about the skill compatible

with the view that thanks to perceptual abilities our veridical experience is object-related and not just appearance-related? The question can be reinterpreted into: is it really necessary to embrace the theory of broad competence in order to account for the broadness of the recognitional act? In this section, I will claim that it is not necessary, and that we can have ability-based disjunctivism based on narrow competence.

The essence of my answer and thereby of my specific disjunctivist theory is a distinction between two types of abilities that we always obtain simultaneously. I want to suggest that there are two types of abilities that could be responsible for shaping perceptual experience: *object-related ability (or O- ability* for short*)*, which relates the agent to external objects, and *appearance-related ability (or OA- ability* for short*)*, which relates the agent to subjective appearances. What puts the agent in contact with external objects in the good case is the *exercise* of an O-ability.[34] What accounts for illusions and hallucination is the *exercise* of OA-ability.

Another crucial distinction that I want to draw is a distinction between having ability and exercise of ability. I will use ability and skill as synonyms. For instance, LP from the above example can plausibly be said to have the skill to recognize Earthly objects, but he is not in a position to exercise that skill. This distinction partly relies on my commitment to the theory of narrow competence. Having the skill is a narrow notion, whereas exercise of skill should be understood as a broad notion. By broad here I mean that the exercise of ability cannot be decomposed into a mere conjunction of internal and external conditions.

I have two reasons to claim that the exercise of ability is broad, but possession of the skill narrow. First, on the view that Sosa, and I have been defending, the exercise of the skill entails: having the skill plus the presence of proper circumstances plus success resulting from that skill. But success resulting from skill cannot be decomposed into having the skill plus the

presence of proper circumstances. In other words, success from skill cannot be decomposed into a mere conjunction of its internal and its external component. It is possible both in trivial and in non-trivial sense to have the internal and the external components satisfied in a conjunction, and yet not to have success from competence. In the trivial sense, this would be a case in which one has a skill, and the appropriate circumstances are present, but one does not try to exercise the skill. In the non-trivial sense this is illustrated by all Gettier cases. For example, take the often quoted story of a boss who enters his office and sees the indistinguishable twin of his secretary, of the existence of whom he is unaware, while his secretary is behind the door in the same room. He forms the belief that his secretary is in the room. The belief is true. It satisfies the external condition, truth, he indeed has the skill to recognize his secretary. Recognitional abilities can be not just of kinds, but also of persons. However the belief of the boss is not true as a result of the exercise of his skill to recognize his secretary. In fact he fails to exercise the skill under question. In this case, he has the skill, and the opportunity to exercise it, i.e. his secretary is there in the room. But he fails to exercise it. So, along the line of my argument in chapter 2 I treat the exercise of the skill (success resulting from skill) as a success notion, which is prime and broad.

I don't further assume that the environment in which the skill is exercised has to be appropriate in Millar's sense. I do not take it that one can exercise a recognitional ability only in environments where in general the appearances of the object of recognition are representative of these objects. I contend that simply hitting the truth as a result of a competence is enough for the exercise of that competence and therefore for knowledge. Recognitional O-ability is strictly speaking object-related. It is not environment related. It is limited to recognizing the object. It does not include also the task of recognizing whether the environment is appropriate, unless of course that is the object targeted for recognition. In the

[34] Sosa (in Bonjour and Sosa 2003, p. 100) distinguishes a special kind of knowledge which he calls "objectual knowledge". Perhaps something like objectual knowledge is the direct result of the exercise of object-related

Barney's case, it is the barn itself that is the targeted object of recognition, not whether appearances of barns in that environment are representative of real barns or not. This is why I claim that Barney, unlike the Gettier subject, can exercise his recognitional ability.

Now in contrast to the exercise of a skill, mere possession of the skill is narrow, and it is not to be described as prime. Having the skill entails having a properly firing neural network in the head in the presence of the object or appearance of the object. The possession of the skill consists of these two modes – it can be related to appearances, or it can be related to objects. Its having these two modes accounts for the possibility to learn a skill in simulational environments. It explains why the astronaut has the ability to survive in the outer space, and why LP and some BIVs can possess the skills to recognize Earthly objects.

Let me give a different example to illustrate more precisely what I mean. Consider the following thought experiment. Suppose that a scientist trains a BIV (including a visual system and eyes) for becoming an embodied perceiver. At certain stage of the BIV's training, the scientist is projecting 2D images on both of BIV's retinas. With every projection, the well-trained BIV makes a top down checking by sending motor signals to its eyes to explore the "environment". Every time the BIV sends these signals, the scientist moves the projected images so as to create a perfect illusion in the BIV as if there are really 3D objects around it. BIV is doing its epistemic duty involved in normal perceiving (making top-down checking), it does everything that is up to it to exercise its seemingly O-ability.

I want to say that this kind of BIV has the ability to perceive, an O-ability. If the scientist had put real objects in front of BIV's eyes, instead of just projected images, it would very probably see real objects. However, in the given circumstances, the BIV cannot exercise its O-ability. On certain occasions, the BIV exercises its OA-ability - the ability to recognize appearances of certain objects as being of these objects. Note that OA-ability does not always

---

ability, but I leave this issue open.

include the possibility to recognize that these are just appearance, and not the objects themselves. In some well-known cases appearances can be indistinguishable from the objects themselves.

I said that I have two reasons to claim that the exercise of O-ability is prime. The second reason is because during the exercise of an O-ability, it is not only that the skill is responsible for cognitively reaching to the object, but also the internal mental state assimilates or absorbs the external objects. The second statement commits me to the credo of metaphysical disjunctivism. I believe that the most specific mental factor of the state of exercise of the O-ability and the state of exercise of OA-ability is different. I think so, because perceptual ability seems to have a strong connection to experience. So, it is difficult to see how perceptual ability can be essentially relational without experience being essentialy relational, since ability in a hinges on experience although it also shapes the experience.. Besides the fact that the most specific mental kind is different, sufficiently many other mental factors can be shared, and this explains why the two states can sometimes be indistinguishable. On its side, indistinguishability can provide a good basis of learning from simulations. I leave the question about which are the shared mental kinds open. I suggest here only that they can have different descriptions: seeing 3D objects, and constructing 3D objects out of 2D ones

5.5 Conclusion

To sum up, the theory suggested and defended in this chapter is that veridical perceptual experience takes in an object in virtue of exercise of perceptual ability. Perceptual ability is a refined skill to recognize objects directly. The process of refinement is underlined and enabled by exercise of propensities. However, I claimed that it is not a result of exercise

of virtues creditable to the agent, because the refinement and its motivation are not consciously accessible by the agent.

Next, I argued that ability that is being exercised in the veridical perceptual case is different from the ability that is being exercised in the bad case. The first one is O-related. The second one is OA-related. They are different because the enabling conditions are different, and because the exercise of O-ability essentially involves an external object. The exercise of O-skill is prime. This is a precondition for having access to the external objects, and therefore perceptual knowledge. Eventually, it is the exercise of O- recognitional ability that links us to the external objects.

Also, I assumed that the most specific mental factor in the veridical case is different from that in the non-veridical case. I think that if we defend an ability-based disjunctivism, it is also fair to embrace some form of metaphysical disjunctivism too, but I leave the issue of its modesty open. At the same time we can allow for significant mental commonalities in the two cases. I did not provide a more specific list of mental commonalities though.

What do we achieve? I suppose we achieve three main things. First, hopefully we block the argument from illusion raised against the traditionalist view by rejecting an essential common factor in the exercise of abilities in the good and in the bad cases. Second, my theory explains how perceptual knowledge can be seen as an achievement, without involving the implausible assumption that it is a virtue creditable to the agent. Third, our theory does not face the indistinguishability problem that is usually raised against metaphysical disjunctivism: when the two different abilities (O- and OA-) are being exercised, the resulting experience of the subject could feel mentally the same, although it is different in nature.

CONCLUSION

Now let us see the dialectics of the separate elements discussed in the chapters. I claimed that knowledge is apt* success or apt* mental state that covers both instances of propositional and practical knowledge. Apt* success is a success from intelligent ability which is an ability refined through intelligent acts, or acts exercised with teleologically appropriate propensities.

Intelligent ability has internal and external bases. The internal basis of the ability has to be safe, i.e. the belief/action must not easily fail to be true/successful if based on that basis. Also, it has to be normalized safe, attuned to a particular range of circumstances. The requirement that the agent has the basis safely is not essential. Hence the definition of knowledge as apt* success covers also cases of knowledge without virtue. The notion of credit needed in epistemology is important to explain the anti-accidentality condition for knowledge. But we have no reasons to think that knowledge can be credited always primarily to the cognizer. In many cases, it is essentially credited to her sub-personal cognitive system. The system does the job of cultivating the anti-accidentality condition – the internal basis of ability – which then enables the agent to grasp the objects of her knowledge directly and justifiably on the conscious level.

The external basis of the ability presupposes various forms of environmental luck. Knowledge is compatible with mitigated bad luck, absence of bad luck, and positive luck. It is not compatible with possible worlds that accommodate the hunches of the cognizer, and in this sense, it is not compatible with the full presence of enabling circumstances. External circumstances have to be normalized safe; they have to be among the range of circumstances to which the ability has been attuned. However, the actual situation of the cognizer need not, by all means, be counterfactually safe. Knowledge is possible even under close risk of failure in cognizer's actual situation.

The relation of aptness, the way in which the internal and the external components have to be connected for knowledge to occur, is claimed to be a prime relation. Only under that condition the definition of knowledge as success from ability is sufficiently immune to counterexamples. The main counterexample against this account is the possibility of deviant causal chains. An apt relation it is not vulnerable to this counterexample, because in cases of deviant causal chains the agent does not appropriately grasp the situation. But intelligent ability is not infallible, so its exercise does not presuppose the notion of knowledge. These considerations are supposed to show that the notion of 'success from ability' explains why knowledge is an analyzable concept. Even if the KSA principle does not provide an entirely non-circular definition of knowledge, this definition is still informative.

Finally, the view has an application in the domain of perception. Involving intelligent abilities in the debate on perception explains how our perceptual experience grounds knowledge about the external objects. Ability that is being exercised in the veridical perceptual case is different from the ability that is being exercised in the non-veridical case. The first one is object–related (an O-ability). The second one is appearance-related (an OA-ability). They are different, because the enabling conditions are different, and also because the exercise of O-ability essentially involves an external object.

What do we achieve as a result of the investigation? We achieve a sustainable definition of knowledge that covers both practical and theoretical knowledge. The requirement of intelligent ability is more minimal, and covers more cases of knowledge than the requirement of virtue. Also, the notion of credit, I have been defending, integrates the anti-luck condition. So, there is no need to postulate an anti-luck condition as a separate condition. In the domain of perception, the theory hopefully blocks the argument from illusion raised against the traditionalist view, by rejecting an essentially common factor between the state of exercising of O-ability, and the state of exercising OA-ability. I am aware that more

considerations need to be given in favor of the purported difference. At the same time the theory resist the indistinguishability problem that is usually raised against the metaphysical disjunctivism: when the two different abilities (O-ability and OA-ability) are being exercised, the resulting experience of the subject could feel mentally the same, although it is arguably different in nature.

BIBLIOGRAPHY

Annas, J. (2003) "The Structure of Virtue" in *Intellectual Virtue: perspectives from ethics and epistemology*, ed. M DePaul and L. Zagzebski, Oxford University Press, pp.15-33.

Aristotle, *Nicomachean Ethics.* Translated by W. D. Ross *The Internet Classics Archive* by Daniel C. Stevenson, Web Atomics. World Wide Web presentation is copyright (C) 1994-2000, http://classics.mit.edu/Aristotle/nicomachaen.html

Armstrong, D. M. (1973) *Belief, Truth and Knowledge*, Cambridge: Cambridge University Press.

Bergmann, M. (1997) "Internalism, Externalism and the No-Defeater Condition" *Synthese* 110 (3): pp. 399-417.

Bernecker, S. (2008). "Agent Reliabilism and the Problem of Clairvoyance" *Philosophy and Phenomenological Research* 76 (1): pp.164–172.

Bernecker.S (2006) "Prospects for Epistemic Compatibilism" *Philosophical Studies* 130 (1): pp.81-104.

Boghossian, P. and Williamson, T. (2003) "Blind Reasoning" *Proceedings of the Aristotelian Society*, Supplementary Volume, 77: pp 225-248

BonJour, L. (1985), *The Structure of Empirical Knowledge*, Cambridge, Massachusetts: Harvard University Press.

Bonjour, L. (1993) "Recent Work on the Internalism–Externalism Controversy" in Dancy, J. and Sosa E. (Eds.) *A Companion to Epistemology*, Blackwell Publishing pp. 33-43.

BonJour, L. and Sosa, E. (2003) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. *A Debate between Laurence BonJour and Ernest Sosa*, in the Series *Great Debates in Philosophy*, Blackwell Publishers.

Breyer, D. and Greco, J. (2008). "Cognitive Integration and the Ownership of Belief: Response to Bernecker" *Philosophy and Phenomenological Research* 76 (1): pp.173–184.

Byrne, A. and Logue, H. (2008) "Either/Or". In A. Haddock & F. Macpherson, (Eds.) *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press.

Byrne, *A. and* Logue, H. (Eds.) (2009) *Disjunctivism: Contemporary Readings*. Cambridge MA: The MIT Press.

Code, L. (1984) "Toward a 'Responsibilist' Epistemology," *Philosophy and Phenomenological Research*, 45 (1): pp.29–50.

Code, L. (1987) *Epistemic Responsibility*, Hanover: University Press of New England and Brown University Press.

Cohen, S. (1984) "Justification and Truth" *Philosophical Studies* 46: pp. 279-96.

Cohen, S. and K. Lehrer. (1983) “Justification, Truth, and Knowledge”. *Synthese* 55: pp.191-207.

Comesana, J. (2002) “The Diagonal and the Demon” *Philosophical Studies* 110: pp.249-66.

Descares, R. *Meditations on First Philosophy*, Translated by John Cottingham, *(*1996) Cambridge: Cambridge University Press.

Descartes, R. *Rules for the Direction of the Natural Intelligence: A Bilingual Edition of the Cartesian Treatise on Method*. Translated and edited by George Heffernan (1998). Amsterdam/Atlanta, Georgia: Rodopi.

Dretske, F. (1971) “Conclusive Reasons” *The Australasian Journal of Philosophy* 49: pp. 1-22.

Dutant, J. (2010) “Two notions of safety” *Swiss Philosophical Preprints* 87.

Foley, R. (1987) *A Theory of Epistemic Rationality* Cambridge: Harvard University Press.

Cooper J.M., Hutchinson D.S. (Eds.) *Plato: Complete* Works, Hackett Publishing Company.

Gettier, Edmund L. (1963) "Is Justified True Belief Knowledge?" *Analysis* 23, pp. 121–123.

Goldman, A. (1967) "A Causal Theory of Knowing" *The Journal of Philosophy* 64: pp. 357-372.

Goldman, A. (1976) “Discrimination and Perceptual Knowledge” *The Journal of Philosophy* 73: pp. 771-91

Goldman, A. (1986) *Epistemology and Cognition*, Cambridge, MA: Harvard University Press

Greco, J. (2000a) *Putting Skeptics in their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*, Cambridge University Press.

Greco, J. (2000b) “Two Kinds of Intellectual Virtue,” *Philosophy and Phenomenological Research* 60: pp. 180-85.

Greco, J. (2001) “Virtues and Rules in Epistemology,” in L. Zagzebski and A. Fairweather, (Eds.) *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Oxford: Oxford University Press.

Greco, J. (2002) “Virtues in Epistemology,” in *The Oxford Handbook of Epistemology*, P. Moser (Ed.) Oxford: Oxford University Press.

Greco, J. (2003), “Further Thoughts on Agent Reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi”, *Philosophy and Phenomenological Research* 66: pp. 466–480.

Greco, J. (2004) “How to Preserve your Virtue while Losing your Perspective” in Greco (Ed) *Sosa and his Critics*, Oxford: Blackwell.

Greco, J. (2005) "Justification Is not Internal." In Steup, M. and Sosa, E. (eds. 2005). *Contemporary Debates in Epistemology*. Oxford: Blackwell. pp. 257 - 269.

Greco, J. (2007). "Discrimination and Testimonial Knowledge" *Episteme* 4 (3): pp. 335-351.

Greco, J. (2009) *Achieving Knowledge*, Cambridge: Cambridge University Press.

Greco, J. (2009) "Knowledge and Success From Ability" *Philosophical Studies* 142 (1).

Greco, J. in (*forthcoming*), "The Value Problem", in *Epistemic Value*, eds. A. Haddock, A. Millar and D. H. Pritchard, Oxford: Oxford University Press.

Greco, J. and Turri, J. (2010) "Virtue Epistemology", *The Stanford Encyclopedia of Philosophy* Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2010/entries/epistemology-virtue/>.

Hinton, J. M. (1973) *Experiences* Oxford: Clarendon Press.

Hookway, C. (1994) "Cognitive Virtues and Epistemic Evaluations," *International Journal of Philosophical Studies*, 2 (2): pp. 211–227.

Jaegwon, K. (1988) "What is 'Naturalized Epistemology'?" *Philosophical Perspectives*, Vol. 2 (1988): pp. 381-405.

Kvanvig, J. (1992) *The Intellectual Virtues and the Life of the Mind*, Savage, Maryland: Rowman and Littlefield.

Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.

Lackey, J. (2009) "Knowledge and credit" *Philosophical Studies*, 142: pp. 27–42.

Lehrer, K. (1990) *Theory of Knowledge*, Westview.

Lehrer, K. and Paxson, T. (1969) "Knowledge: Undefeated Justified True Belief" *The Journal of Philosophy* Vol. 66 (8): pp. 225–237.

Lepock, C. (2006) "Adaptability and Perspective." *Philosophical Studies* 129 (2): pp. 377-391.

Lewis, C. (1895) "What the Tortoise Said to Achilles" *Mind*, 4: pp. 278–80.

Martin, M. G. F. (2004) "The Limits of Self-Awareness" *Philosophical Studies*, 120: pp.37–89.

Mathen, M. (2005) *Seeing, Doing, and Knowing: A Philosophical Theory of Sense-Perception*, Oxford: Clarendon Press.

McDowell, J. (1982) “Criteria, Defeasibility, and Knowledge” *Proceedings of the British Academy*, 68: 455–79. reprinted in (1998) *Meaning, Knowledge, and Reality*, Cambridge, MA: Harvard University Press: pp.369–94.

Millar, A. (2008 a) “Perceptual-Recognitional Abilities and Perceptual Knowledge” in Adrian Haddock and Fiona Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge* (Oxford: Oxford University Press).

Millar, A. (2008 b) “Disjunctivism and Skepticism” in John Greco (Ed.) *The Oxford Handbook of Skepticism* (New York: Oxford University Press: pp. 581-604.

Millikan, R.G. (2000) *On Clear and Confused Ideas: An Essay about Substance Concepts* Cambridge: Cambridge University Press.

Moffett, M. and Bengson, J. (2007) "Know How and Concept Possession" *Philosophical Studies*, 136: pp. 31–57.

Montmarquet, J. (1993) *Epistemic Virtue and Doxastic Responsibility*, Lanham: Rowman and Littlefield.

Neipher, S. (2008) “A Virtue Account of Perception” in Neipher, S. (Ed.) *Virtue Epistemology: Moivation and Knowledge*, Continuum: pp. 41-59.

Norman, D. (1993) “Cognition in the Head and in the World: An Introduction to the Special Issue on Situated Action” *Cognitive Science*, Special Issue: Situated Action 17: pp.1-6.

Pappas, G. "Internalist vs. Externalist Conceptions of Epistemic Justification", *The Stanford Encyclopedia of Philosophy (Fall 2008 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/justep-intext/>.

Peacocke, C. (1999) *Being Known*, Oxford University Press.

Plato *Complete Works by Plato* John M. Cooper and D. S. Hutchinson (Eds.) (1997), Hackett Publishing Company.

Plantinga, A. (1993). *Warrant: The Current Debate*, Oxford: Oxford University Press.

Price, H. H., 1950. *Perception*, London: Methuen.

Prichard, H. A. (1950) *Knowledge and Perception*, Oxford: Clarendon Press.

Prinz, J. (2006) “Beyond Appearances: The Content of Sensation and Perception” in *Perceptual Experience* Ed. T. Szabo Gendler and John Hawthorne: pp. 434-61.

Pritchard, D. (2007) “Recent Work on Epistemic Value,” *American Philosophical Quarterly*, 44: pp. 85–110.

Pritchard, D., Millar, A. and Haddock, A. (2010) *The Nature and Value of Knowledge: Three Investigations*, Oxford University Press.

Quine, W.V.O. (1969) *Ontological Relativity and Other Essays*, New York: Columbia University Press.

Riggs, W. (2003) "Understanding 'Virtue' and the Virtue of Understanding," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. Michael DePaul and Linda Zagzebski, Oxford: Oxford University Press.

Riggs, W. (2007) "Why Epistemologists Are So Down on Their Luck," *Synthese*, 158: 329–344.

Riggs, W. (2009) "Two Problems o f Easy Credit," *Synthese* 169: pp. 201–216.

Roberts, R.C. and Wood W.J. (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*, Oxford University Press.

Robinson, H. (1985). "The General Form of the Argument for Berkeleian Idealism", in J. Foster and H. Robinson (eds.), *Essays on Berkeley*. Oxford: Oxford University Press: pp.163-86.

Rumfitt, I. (2003) "Savoir Faire" in *The Journal of Philosophy*, 100(3): pp.158-166.

Ryle, G. (1949) *The Concept of Mind*, London: Hutchinson's University Library.

Slater A. (Ed.) (1998) *Perceptual Development: Visual, Auditory, and Speech Perception in Infants* Psychology Press, UK.

Snowdon, P. (1980–1) "Perception, Vision, and Causation", *Proceedings of the Aristotelian Society*, 81: pp. 175–92.

Snowdon, P. (1992) "How to interpret 'direct perception'" in T. Crane (ed.) *The Contents of Experience: Essays on Perception* (Cambridge: Cambridge University Press), pp. 49-78.

Snowdon, P. (2005) "The Formulation of Disjunctivism: A Response to Fish" in *Proceedings of the Aristotelian Society* 105: 129-41.

Sosa, E. (1980) "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge" *Midwest Studies in Philosophy* Vol. V: 3-25.

Sosa, E. (1991) *Knowledge in Perspective*, Cambridge: Cambridge University Press.

Sosa, E. (1993) "Abilities, Concepts, and Externalism" in J. Heil and A. Mele (Eds.) *Mental Causation*. Oxford: Clarendon Press.

Sosa, E. (1994) "Philosophical Scepticism and Epistemic Circularity" *Proceedings of the Aristotelian Society*: pp. 263-90.

Sosa, E. (2007) *A Virtue Epistemology: Apt Belief and Reflective Knowledge. Volume One*, Oxford University Press.

Stanley J. and Williamson T. (2001) "Knowing How" in *Journal of Philosophy* 98 (8): pp. 411-44.

Steup, M. (2004). Internalist Reliabilism. *Philosophical Issues* 14 (1): pp. 403–425

Steup, M. (2008) "The Analysis of Knowledge", *The Stanford Encyclopedia of Philosophy)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2008/entries/knowledge-analysis/>.

Swain, M. (1974) "Epistemic Defeasibility" *The American Philosophical Quarterly*, 11(1): pp. 15-25.

Treffert D.A. (2009) "The savant syndrome: an extraordinary condition. A synopsis: past, present, future" in *Philosophical Transactions of the Royal Society of Biological Sciences*, Vol. 364 no. 1522: pp. 1351-1357.

Van Cleve, J. (1979). Foundationalism, Epistemic Principles, and the Cartesian Circle. *Philosophical Review* 88 (1): pp. 55-91.

Van Fraassen, B. (1980) *The Scientific Image*, Oxford University Press.

Vogel, J. (2008). "Epistemic Bootstrapping." *Journal of Philosophy* 105 (9): pp. 518–539.

Williamson, T. (2000) *Knowledge and Its Limits*, Oxford University Press.

Williamson, T. (2004) "Sosa on Abilities, Concepts, and Externalism" in Greco (Ed.) *Sosa and his Critics*, Oxford: Blackwell.

Zagzebski, L. (1996) *Virtues of the Mind*, Cambridge: Cambridge University Press.

Zagzebski, L. (1999) "What is Knowledge?" in *The Blackwell Guide to Epistemology*, J. Greco and E. Sosa (Eds.) Malden, MA: Blackwell.

Zagzebski, L. (2000) "Responses," *Philosophy and Phenomenological Research* 60: pp. 209-221.