

© 2014

**Lisa Miracchi**

**ALL RIGHTS RESERVED**

# GETTING THINGS DONE

BY LISA MIRACCHI

A dissertation submitted to the  
Graduate School—New Brunswick  
Rutgers, The State University of New Jersey  
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

Written under the direction of

Ernest Sosa and Brian McLaughlin

and approved by

---

---

---

---

New Brunswick, New Jersey

May, 2014

## ABSTRACT OF THE DISSERTATION

### Getting Things Done

by Lisa Miracchi

Dissertation Directors: Ernest Sosa and Brian McLaughlin

In *Getting Things Done*, I develop and defend a new theory of achievement. An achievement, as I use the term, is the crown performance of a cognitive or practical domain. In the perceptual domain, the achievement is perceiving things as they are; in the epistemic domain, it is knowing that *p*; in the practical domain, it is intentionally accomplishing what one intends. In each of these domains there are corresponding cases of failure that deviate from achievements in interesting ways, and that also deserve explanation. These include hallucinations, perceptual illusions, cases of justified false belief, and cases where we fail to do what we try to do. Traditionally, theorists have supposed that achievements and their corresponding failures may both be explained in terms of *neutral performances*—mental states or actions that may obtain both in cases of achievement and in cases of failure—and *non-personal* (non-mental, non-agential) conditions. For example, visually representing, believing, and intending have all been posited as neutral common factors that help to explain the achievements and corresponding failures in their respective domains. In my dissertation, I argue against this approach and develop an *achievement-first* alternative.

In “Achievements and Exercises”, I argue against the traditional common factor approach to competences, and propose an achievement-first theory—the *dual exercise*

*account*. According to it, there are no neutral exercises of competence; exercises are either constitutively achievements or constitutively failures. In “Competence to Know”, I apply the dual exercise account to the epistemic domain, and propose a *direct virtue epistemology*, on which knowledge is a manifestation of a competence to know, not to believe truly. In “A Virtue Aistheology”, I apply the dual exercise account to the perceptual domain.

In these three chapters, I provide a new framework for theorizing about mental phenomena. I show how an achievement-first virtue-theoretic approach can be an explanatory rival of the traditional common factor approach. It makes progress both in answering questions that arise on any achievement-first approach, such as what it is for achievements to be mental states in their own right, as well as traditional problems, such as how achievements are related to corresponding failures.

## Acknowledgements

It really does take a village.

First, thanks to my dissertation co-advisors Ernest Sosa and Brian McLaughlin, to the rest of my dissertation committee, Susanna Schellenberg and Jason Stanley.

Thanks also to friends, colleagues, and teachers who have given invaluable comments on this work: Andy Egan, Ben Bronner, Blake Roeber, Bob Beddor, Carlotta Pavese, Dan Harris, Daniel Singer, David Black, E.J. Green, Eliot Michaelson, Frances Egan, Georgi Gardiner, Herman Cappelen, Jennifer Nagel, Jessica Brown, John Turri, Jonathan Ichikawa Jenkins, Kate Nolfi, Katherine van Elswyk, Kurt Sylvan, Lucy Jordan, Matt Benton, Megan Feeney, Nico Kirk-Giannini, Stephanie Leary, Pamela Robinson, Peter Klein, Peter van Elswyk, Rachel McKinney, Ricardo Mena, Rodrigo Borges, Stephanie Leary, Susanna Schellenberg, Ted Poston, Will Fleisher, Zoe Jenkin, and all the other people I'm sure I've forgotten to mention here.

For their invaluable support, thanks to Ricardo Mena Gallardo, Janice Miracchi, Silvano Miracchi, David Miracchi, John Sheffield, Robert Sinnott, and my dear canine companion Owen.

Lastly, thanks to Jennifer Malin (December 21, 1986–October 26, 2013), who has been, and always will be, an inspiration.

## **Dedication**

For Jennifer Malin, who will never be able to write her dissertation.

# Table of Contents

<b>Abstract</b> . . . . .	ii
<b>Acknowledgements</b> . . . . .	iv
<b>Dedication</b> . . . . .	v
<b>1. Introduction</b> . . . . .	1
<b>2. Achievements and Exercises</b> . . . . .	11
2.1. What Do We Want from a Theory of Competence? . . . . .	15
2.1.1. Chomsky’s Competence/ Performance Distinction . . . . .	15
2.1.2. Which Are the Competent Performances? . . . . .	18
2.1.3. Desiderata for a Philosophical Theory of Competence . . . . .	19
2.2. Comparison of Existing Approaches . . . . .	21
2.2.1. The Success Approach . . . . .	21
Deviant Causal Chains . . . . .	23
2.2.2. The Achievement Approach . . . . .	28
2.3. Fallible Competences, Factive Exercises . . . . .	30
2.3.1. A Middle Way? . . . . .	30
2.3.2. The Modal Argument Against Fallible Competences with Factive Exercises . . . . .	31
2.3.3. The Dual Exercise Account . . . . .	33
2.4. Conclusion . . . . .	44
<b>3. Competence to Know</b> . . . . .	47
3.1. Indirect Virtue Epistemology . . . . .	49
3.1.1. Between a Rock and a Hard Place . . . . .	51

3.1.2.	Sosa (2007)	52
3.1.3.	Greco (2012)	54
3.1.4.	Sosa (2010)	55
3.1.5.	A General Procedure	58
3.1.6.	Safety to the Rescue?	61
3.2.	Direct Virtue Epistemology	62
3.2.1.	Knowledge Is an Exercise of a Competence to Know	62
3.2.2.	Knowledge Is a Mental State	64
3.2.3.	Knowledge and Justification	66
3.2.4.	Mere Belief	74
3.3.	Conclusion: What Kind of Theory is Direct Virtue Epistemology?	76
3.3.1.	The Nature and Normativity of Knowledge	77
3.3.2.	Informative Partial Metaphysical Analysis of Knowledge	78
3.3.3.	Full Metaphysical Analysis of Justification	80
3.3.4.	Moving on	80
<b>4.</b>	<b>A Virtue Aistheology</b>	<b>82</b>
4.1.	Berkeley's Puzzle and Adequate Explanatory Scope	84
4.1.1.	Is Campbell's Relationalism an Equally Adequate Alternative?	88
4.2.	The Perceptual Relation and Scientific Investigation	90
4.2.1.	The Perceptual Relation	91
4.2.2.	Scientific Investigation of Competent Performance	92
4.3.	The Rational Role of Perception	97
4.4.	Hallucinations, Illusions, and Subjective Indiscriminability	102
4.4.1.	Perceptual Competences	102
4.4.2.	Hallucinations	106
4.4.3.	Illusions	108
4.4.4.	Perceptual Experience	109
4.5.	Conclusion	110



## Chapter 1

### Introduction

Fred Dretske opens his 1986 paper “Misrepresentation” with the following two sentences:

Epistemology is concerned with knowledge: how do we manage to get things right? There is a deeper question: how do we manage to get things wrong?

—Dretske (1986)

This line of thinking is a unifying theme of much of 20th century philosophy of mind and epistemology, which as a result is seen as a separate discipline. The idea is that the *core*, most fundamental mental or person-level phenomena are phenomena that can get it *wrong* about how things are, or can *fail* to bring about what we aim for. They do not entail success or achievement—they are perceptual experiences, beliefs, desires, and intentions, *not* perceivings, knowings, and doings.

This way of thinking about mental states and actions excises much more than knowledge from the realm of the mental. Jerry Fodor, for example, is very clear about how his conception of the mind, and the research program for psychology that it generates, is incompatible with thinking not only of knowledge but any factive states as mental states:

Take, for example, knowing that such-and-such, and assume that you can't know what's not the case. Since, on that assumption, knowledge is involved with truth, and since truth is a semantic notion, it's going to follow that there can't be a psychology of *knowledge*... . Similarly, it's a way of making a point of Ryle's to say that, strictly speaking, there can't be a psychology of perception if the formality condition is to be complied with. Seeing is an achievement; you can't see what's not there.

—Fodor (1980), pp. 652-653.

This idea articulated by Dretske and Fodor that the *real* mental states—i.e. the ones that are the object of science and which do the serious philosophical explanatory work—are non-factive was dogma in modern and contemporary philosophy until very recently.

There are several explanations one might give of this expulsion of perception, knowledge, and other similar mental states from the realm of the mental, but I think that the simplest and most plausible story is that we seem to get so much *explanatory mileage* out of supposing that it is non-factive mental states such as perceptual experiences and beliefs that are metaphysically fundamental. If we posit, say, perceptual experience as the most fundamental state in the perceptual domain, then it seems that we can explain perceiving things as they are as well as hallucinations and illusions in terms of this common state. We can also make both philosophical and scientific progress in understanding *what it is* to have a perceptual experience and *in virtue of what* we have the experiences that we do. This project, then, allows us to have a unified account of the perceptual, and to bring it all within the domain of science.

Likewise for knowledge. We begin by supposing that beliefs and perceptual experiences are the core mental phenomena, and then explain what it is to have a *justified* belief either in terms of non-mental relations to the world (traditional externalism) or in terms of other non-factive mental states (traditional internalism). We then explain what it is to have knowledge in terms of further non-mental relations. If this project can be carried out, then we only need to naturalize perceptual experiences and beliefs. We have a unified, wide-reaching, naturalistic, explanatory theory.

The trouble, of course, is that these theories have never been as clean as they have been made out to be. Even explaining justified belief in terms of belief and its relations to other mental states or non-mental relations to the environment has been exceedingly difficult. But explaining perception and knowledge in terms of perceptual experiences, beliefs, and non-mental relations has presented a distinctive problem. Perceiving and knowing don't just seem to be a matter of *getting things right*, as Dretske supposes. Rather, Fodor was closer; they seem to be achievements, a matter of *getting things done*.

Cases that show the distinction between *getting things right* and *getting things done* have been around for a good while, and they have supported the idea that perceiving, knowing, and other similar states are a matter of getting things done, rather than merely getting things right. For example, consider the following case:

**Veridical Hallucination.** Sandra's new puppy has been giving her a headache. She reaches for pain killers but her friend Zoe, knowing how trying it is to have a new puppy, decides to play a trick on her, replacing her pain killers with pills that make one hallucinate black and white dogs. Sandra takes a pill, which causes her to have a perceptual experience as of a black and white dog on the floor before her. In fact, there is a black and white dog on the floor before her. (Her puppy is black and white).

This is a case of what some people call "veridical hallucination"—Sandra has a perceptual experience that "matches" a real-world scenario, but she hallucinates (purportedly) because the experience is not appropriately related to the fact that makes her experience veridical. She gets things *right*, but she doesn't perceive. Perceiving thus cannot merely be a matter of getting things right.

Edmund Gettier (1963) developed similar cases for knowledge, and they have received considerably more attention than their counterparts for perception. Knowledge is not just justified belief that gets things right—i.e. justified belief that is true—it is a matter of *achieving* the truth.

One might try, nevertheless, to explain what it is to get things done in terms of what it is to get things right and further non-mental conditions. Indeed, the last 50 years of work in epistemology have largely been shaped by attempts to solve the Gettier Problem in this way—that is, by attempts to show what non-mental or non-factive mental additions to having a justified true belief are necessary for, and make it the case that, one knows.

This project has been an utter failure. It has been such a failure that the attempt to explain knowledge in terms of justified belief has largely been abandoned. Following Williamson (2000), many have taken the lesson to be that we should put knowledge *first*. Knowledge is a mental state *in its own right*—that is, not merely in virtue of being partially constituted by belief or justified belief.

[T]he claim that knowing is a state of mind is to be understood as the claim that there is a mental state being in which is necessary *and sufficient* for knowing *p*. In short, knowing is *merely* a state of mind. This claim may be unexpected. On the standard view, believing is merely a state of mind but knowing is not, because it is factive: truth is a non-mental component of knowing.

—Williamson (2000), p. 22

However, many have also *not* taken up this banner. They continue to believe that knowledge is not properly mental, that it derives its status as mental from being partially constituted by justified belief. We may not be able to explain exactly what extra is required, the thinking goes, but there must be an answer to this question.

Why has there been such resistance to the idea that knowledge is a mental state in its own right? The answer, I think, is again quite simple: the knowledge-first program is, to date, *less explanatory* than the common factor program. We do not have a great sense of how knowledge *could be* a mental state. What kinds of things are mental states, if knowledge is one? How might such mental states obtain in virtue of non-mental facts? We have to revise the mainstream story so considerably that it is difficult to understand exactly what the knowledge-first position is.

We also don't have a great sense of how knowledge relates to *failures*, such as justified false beliefs and Gettier cases, that seem to be closely related to them; nor do we understand why knowledge and its corresponding failures seem to share mental states—belief and justified belief—and to have normative statuses seemingly in virtue of what they have in common. All of these mysteries generated by the knowledge-first research program have left many unmoved by Gettier-related considerations. The common factor program, many think, is still our best shot despite its problems. It is just not clear how a knowledge-first program could be equally explanatory.

The same story can be told for why perception-first theories have not caught on, and here things are even worse for perception-firsters. In philosophy of perception veridical hallucinations have not been the main motivator of perception-first theories, but rather concerns about the subject's *perspective*—questions about how the traditional model can explain the kind of first-person awareness we have of our surroundings, how perception can make yet reasonable *to the subject* that she believe that things

are a certain way. While I am very sympathetic to these worries, they are much more difficult to press than the clear counterexamples to traditional analyses of perception and knowledge in terms of non-factive mental states and non-mental conditions.

Consider, for example, Susanna Siegel's review (2004b) of John Campbell's *Reference and Consciousness* (2002b), in which she discusses this kind of perception-first objection to the common-factor approach:

Do nonrelational [common factor] views really disrespect the fact (assuming it is one) that seeing the cup enables you to know which particular thing you are talking about when you say "that cup," pointing to a cup before you? ... Campbell criticizes this version of the view on similar grounds: "all that is within the perceiver's subjective life is the demonstrative element itself," and so an experience with such content "cannot by itself, therefore, distinguish between presentation of one object and presentation of another. ... It is, therefore, opaque how the demonstrative element could provide the subject with an understanding of the demonstrative term." (125)

A proponent of ... nonrelational views could agree that the repeatable part of the perception considered apart from external conditions does not "provide [the subject] with knowledge" of which cup is in question, but hold that the repeatable part when had under the right external conditions does. ... It is open to the proponents of the nonrelational views mentioned here to hold that whether experience gives you knowledge of reference depends not just on its phenomenal character or content, but on the external circumstances in which the experience is had. This move deserves discussion. It isn't hugely theory-laden; in fact, it is a starting point in debates in epistemology: there is a strong intuition that you don't know there's a barn in front of you even when you're seeing one if you're in fake barn country. But the move is ignored. Since it seems compatible with Campbell's explanatory constraint, that constraint seems less powerful than Campbell supposes.

—Siegel (2004b), pp. 430-431

Here Siegel refuses to accept Campbell's desideratum, that a theory of perception must explain how perception provides one with a first-person perspective on the object, the kind of perspective that can provide one with a distinctively mental kind of understanding of the reference of a demonstrative. She denies that the sense in which we have understanding or knowledge of the reference of a demonstrative *is* distinctively first-personal, claiming that it is perfectly acceptable to understand the notions of "understanding" and "knowledge" in terms friendly to experience-firsters.

While I side with Campbell here—I agree that knowledge of the reference of a demonstrative is a more first-personal phenomenon than common factor views can explain, there isn't too much to say to Siegel. These intuitions are not as clear-cut as intuitions about Gettier cases, and I agree that these considerations are not sufficient to require abandoning a theory that is otherwise highly unifying and explanatory.

Interestingly for this discussion, Siegel also allies the project of experience-firsters with the belief-first project in epistemology. She is exactly right to do so. The widespread experience-first strategy for perception and the belief-first strategy in epistemology garner support from one another. They are both species of the common factor approach, on which the fundamental mental phenomena are shared in cases of achievement and corresponding cases of failure, their differences explicable in non-mental terms. However, Siegel is not right to say that the idea in epistemology that “external” considerations make a difference to whether or not one knows “isn't hugely theory laden”. It is, but it is just laden with a very widely accepted theory.

What we have are two different ways of thinking about the mental. There are those who put achievements—perceiving, knowing, and so on—*first*, at the core of mental phenomena. Then there are those who put non-factive “common factor” mental states—perceptually experiencing, believing, and so on—*first*. To date, common factor theorists are winning. Their project is more unified and has wider and deeper explanatory scope. It connects up systematically with cognitive science, furthering the naturalistic goals of many philosophers. But I think it is wrong. I think that achievement-first philosophy of mind is the right approach. The deeper question is how we get things done, not how we get things wrong (or right).

The primary aim of this dissertation is not to defend this claim—although I do make some arguments in favor of it—but is instead to develop an achievement-first program can be just as explanatory as the common factor program. Once we have a real achievement-first alternative, I think the view that perception, knowledge, and other achievements are mental in their own right and fundamental to the realm of the mental will easily become mainstream.

In order to accomplish this task, we need an achievement-first tool-kit. We need some tools for understanding what it means to say that achievements are mental in their own right, for understanding the relationship between achievements and the non-mental biological and environmental facts that give rise to them, and for understanding the relationship between achievements and their corresponding failures, including what the two have in common. The tool kit I propose is a theory of competences.

The central ideas are as follows. First, achievements—perceiving, knowing, and so on—are manifestations of competences. Manifestations are the characteristic exercises of competences (but need not be their only exercises). The competences which they manifest are competences to do that very thing. So, for example, a case of knowing on the basis of testimony is a manifestation of a competence to know on the basis of testimony. This characterization of the *nature* of achievements clarifies the claim that achievements are mental states in their own right. It is widely held that exercises of competences (capacities/ abilities/ methods) are mental (or at least person-level actions). To claim that achievements are exercises of competences relates them to a kind that we have at least some independent grasp of.

If achievements are manifestations of competences to achieve, clearly competences cannot be intended to figure as part of a reductive analysis of achievement. Rather they play an analogous role to that of representation in a common factor theory of perception. Typically, common factor theorists about perception hold that the core perceptual phenomenon is the having of a perceptual experience. They then claim that *what it is* to have a perceptual experience is to perceptually represent that *p*, for example that there is a black and white dog before one. This is a content-laden characterization of perceptual experience, and so is not appropriate to serve as-is in a metaphysical analysis of perceptual experience.

However, the characterization serves two very important purposes. First, it facilitates the use of perceptual experiences in explanations of other perceptual phenomena. If what it is to have a perceptual experience is to perceptually represent that *p*, then the theory can suppose that what it is to perceive things as they are is to *veridically*

perceptually represent that  $p$  in the right conditions.

Second, it suggests strategies for its own metaphysical analysis in terms of non-mental phenomena. If what it is to have a perceptual experience as of an  $F$  is to perceptually represent that there is an  $F$ , then perhaps we can break down the problem metaphysically analyzing perceptual experience into two sub-problems: that of analyzing what it is for a subject to have an experience, and that of analyzing what it is for an experience to have representational content.<sup>1</sup>

Providing a metaphysical analysis of what it is for an experience to have representational content has seemed to many a *much* easier task than that of providing an analysis of what it is for a subject to have an experience as of something's being a certain way. If the problem can be productively approached this way, so much the better for the common factor theory. (Now, I think this divide-and-conquer strategy doesn't work. I do, however, agree that it would be beautiful if it did, and that this sort of division of problems into sub-problems is exactly the kind of thing that a characterization of the core mental phenomenon should suggest.)

I show how the achievement-firster can appeal to competences to do both explanatory tasks. The claim that *what it is* to be an achievement is to be a manifestation of a competence elucidates the phenomenon in a way that (i) allows for explanation of related phenomena in terms of it, and (ii) suggests strategies for its own metaphysical analysis in terms of non-mental phenomena.

Three self-standing chapters follow this introduction. In chapter 2, "Achievements and Exercises", I argue against the traditional common factor approach to competences, and propose an achievement-first theory: the *dual exercise account*. According to the dual exercise account, there are no neutral exercises of competence; exercises are either constitutively achievements or constitutively failures. I specify some necessary conditions on competence possession and manifestation, and show how those conditions may be used to provide necessary *and sufficient* conditions for the exercises of competence that are constitutively failures, as well as the general notion of *exercise*

---

<sup>1</sup>See Field (1978) for a particularly clear articulation of this strategy.



*of competence*. I also show how the theory connects up in a natural way with current projects in cognitive science, suggesting opportunities for fruitful, systematic investigation of achievements and competences to achieve.

In chapter 3, “Competence to Know”, I argue, contra several recent proposals, that no traditional virtue epistemology can solve the Gettier Problem. I then apply the dual exercise account of competences to the epistemic domain, and propose a *direct virtue epistemology*, on which knowledge is a manifestation of a competence to know, not to believe truly. Direct virtue epistemology is a substantive knowledge-first theory—one that helps to explain the facts in virtue of which a subject knows that  $p$ , and that can account for justification and belief in terms of knowledge and the competences that explain knowledge.

In chapter 4, “A Virtue Aistheology”, I apply the dual exercise account of competences to the perceptual domain, developing and defending my *direct virtue aistheology* (study of perception). According to direct virtue aistheology, seeing things as they are is a manifestation of a competence to see things. I show how the account is an explanatory rival to the widespread common factor theory, while avoiding some of its difficulties. I argue that it can explain (i) how how causal facts may make it the case that the perceiver has experiences that purport to be of her immediate environment without falling prey to Berkeleyan skepticism; (ii) how perception may be nonconceptual and nevertheless rationalize beliefs; and (iii) how illusions and hallucinations may be subjectively indiscriminable from cases where we see things as they are.

In these three chapters, I provide a new framework for theorizing about mental phenomena. I show how an achievement-first virtue-theoretic approach can be an explanatory rival of the traditional common factor approach. It makes progress both in answering questions that arise on any achievement-first approach, such as what it is for achievements to be mental states in their own right, as well as traditional problems, such as how achievements are related to corresponding failures. I show how the account presents new avenues for scientific research, and I apply the account to develop theories of perception and knowledge which provide straightforward and plausible solutions to central traditional problems in the perceptual and epistemic

domains.

## Chapter 2

### Achievements and Exercises

#### Introduction: Why Care About Competences?

We perceive, learn, walk, think, speak, drive, understand what is said to us, pour coffee, infer, and so on. We are really good at these things. We do them all the time, and rarely make mistakes. We are very *competent* at what we do.

While there are clearly important differences between the kinds of competent performances mentioned above—they range from perception to thought to language to action—a general theory of competence would provide a common framework for investigating them and these various domains. We would be able to better understand each of them by understanding their commonalities, and we would also be able to better understand their differences as variations or modifications within this general framework. Thus although clearly the phenomena of interest here are quite diverse, the prospect of a unified account is attractive.

There is, however, an even more compelling reason to be interested in the subject. It seems not to be a mere *coincidence* that we are incredibly competent at what we do, or that the more basic or central the performance to our everyday lives the better we are at it.<sup>1</sup> Perhaps the most fundamental performances are essentially competent performances; perhaps *what it is* to perceive things as they are, make a valid inference,

---

<sup>1</sup>Here and throughout, I will use the term “performance” as a general term to cover all of the states and activities of interest to us. Performances are the states and activities that have necessarily subjects or agents, are paradigmatically or often conscious, and are the typical subjects of philosophy of mind, language, and action, epistemology, and cognitive science. We need a general term for these various human and animal states and activities that are the objects of study in philosophy and cognitive science. This is the best one I’ve got so far. I do not intend to prejudge the whether there is special unity to this class beyond the commonalities mentioned above.

assert that p, and so on, is to manifest one's competence to do things of that sort.<sup>2</sup>

COMPETENCE-FUNDAMENTAL PERFORMANCE THESIS: The fundamental performances are essentially competent performances.

In order to see whether this idea is even initially plausible, we must get clearer on what competent performances are. I will suppose that they have two central features. First, a competent performance *manifests* or is an *exercise of*, a competence seated in the subject. So, for example, one might have a competence to tell apart English Springer Spaniels from Cocker Spaniels by sight. Identifying a dog as an English Springer Spaniel in the appropriate way manifests this competence. Of course this relation between competent performance and competence deserves further elucidation, but hopefully the reader has some prior grasp of the idea.

Second, a competent performance is likely to be successful.<sup>3</sup> This need not be true of exercises of abilities or capacities. So, for example, the novice may be *able* to correctly identify which is the English Springer Spaniel with an English Springer Spaniel and a Cocker Spaniel. Nevertheless, she cannot yet reliably tell English Springer Spaniels apart from Cocker Spaniels, and so neither possesses nor exercises the relevant competence.

Once these two aspects of competent performance are articulated, the thesis becomes quite interesting. The idea that one or both of these features are essential to fundamental performances like perceptual and intentional states is shared by a good deal of naturalistic philosophy of mind and language. For example, Burge (2010) claims that what it is to be a perceptual representational state is to be an exercise (realization) of a capacity to veridically represent:

---

<sup>2</sup>Of course sometimes we do things at which we are not competent or skillful. For example in learning how to tell apart English Springer Spaniels from Cocker Spaniels by sight, it might take some practice. Only after such practice are your identifications competent or skillful. In the beginning, you aren't competently identifying the dogs' breeds at all. The Competence-Fundamental Performance Thesis claims that performances like the novice's identifications of English Springer Spaniels are performances only because they are partially constituted by, or otherwise metaphysically dependent upon, other performances that are competent—for example, thinking about English Springer Spaniels, competent related perceptual identifications, and so on. The fundamental performances are competent performances, and all non-competent or incompetent performances metaphysically depend on them.

<sup>3</sup>The degree of reliability may of course vary with domain, context, etc.

Perceptual states are realizations of individuals' capacities. I think that this claim is apriori.

—Burge (2010), p. 369

Veridicality conditions—conditions for accuracy—partly constitute what the perceptual state is. ... Perceptual accuracy is success in fulfilling the state's veridicality conditions, and its representational function. It is perceiving accurately.

—Burge (2010), p. 379

According to Burge, perceptual states are, as a matter of the kinds of states they are, exercises of capacities of the subject to veridically perceptually represent.

As another example, consider Fodor's asymmetrical causal dependence account of mental representation, on which a symbol represents the presence of a certain feature only if it tends to be tokened by that feature.<sup>4</sup> Only because a certain mental symbol tends to be tokened by cows, say, does it represent cows at all. But then the very facts that make the symbol COW represent the presence of a cow also make it the case that tokens of COW *reliably* veridically represent the presence of a cow. That is, according to this kind of approach, whatever the facts are that determine representation *simpliciter* also determine that the representation is reliably veridical.<sup>5</sup> On this kind of approach, causal and other regularities between cognitive systems and environments not only explain how such cognitive symbols reliably represent features of the environment, but they also explain *what it is* for such symbols to represent in the first place.

These ideas also inform cognitive scientific investigation into intelligent performance. Indeed, they are so widespread that they typically serve as background assumptions to scientific investigation. Here, however, is one example where they are expressed:

---

<sup>4</sup>Fodor (1987, 1990b). Note that here the "only if" flags that reliability is a mere necessary condition. Fodor's theory imposes other conditions that are irrelevant to the present discussion.

<sup>5</sup>See also Putnam (1999). This idea can come apart from another popular view, namely that teleological function determines representational content (e.g., Millikan (1989); Dretske (1986); Burge (2010)). However, this latter view endorses the first idea, namely that what it is to be a representation is to be an exercise of a capacity, in this case a capacity to perform a certain biological function.

“Most of us take completely for granted our ability to see the world around us. How we do it seems no great mystery: We just open our eyes and look! When we do, we perceive a complex array of meaningful objects located in three-dimensional space. ... Yet, when viewed critically as an ability that must be explained, visual perception is so incredibly complex that it seems almost a miracle that we can do it at all. ... How, then, are we able so quickly and effortlessly to perceive the meaningful, coherent, three-dimensional scene that we obviously do experience ... ?

—Palmer (1999), p. 4

Palmer wants to understand what it is about us that makes it so that on particular occasions we quickly, effortlessly, (reliably) perceive the world as it is, and he assumes that this will help us to understand *what it is* to see.<sup>6</sup>

That these ideas are so central to philosophical and scientific thinking about fundamental performances such as vision and thought provides compelling reason to investigate competent performance in general. If we can get clearer on what competent performances and the competences they manifest are, then we will hopefully be able to better understand all of the performances of interest to philosophy and cognitive science, and of the facts in virtue of which they obtain.<sup>7</sup> At the very least, we will be able to better evaluate the Competence-Fundamental Performance Thesis.

So, although my aim here is to develop a theory of competence and not to defend the Competence-Fundamental Performance Thesis, interest in this thesis will help to guide the discussion. In section 1, I clarify the sense of competence that is of interest here and set out three desiderata for a theory of competence. I clarify the purpose of such a theory both for “solely” philosophical issues and as a framework for cognitive scientific inquiry. In section 2, I discuss two strategies that have been pursued in explanations of competent performance, and argue that each has important defects. In section 3, I develop my preferred account of competence, which I call the *dual exercise account*. I show how the theory avoids the problems raised for existing views, as well as how it satisfies the desiderata laid out in section 1.

---

<sup>6</sup>Palmer (1999), pp. 4-5.

<sup>7</sup>For this reason, I will approach the idea of a performance as competent performance directly, not through representation or some other prior mental notion. If competence is really as central as I think it is then we should not need to understand it in terms of any other mental notions.

## 2.1 What Do We Want from a Theory of Competence?

There are many things one might want from a theory of competence, and differences in explanatory interest will lead to different kinds of theories. Because my goal is ultimately to understand intelligent performance, I want a theory of competence that aims to explain competent performance. Such a theory would allow me to investigate the relationship between competent performance and performance in general.

But one might have different goals. One might aim rather to systematize the rules that govern certain aspects of a domain of competent performance, without much attention to the actual performances themselves. This difference in explanatory goals will lead to a very different kind of theory. In order to clarify this difference, I will take a moment to discuss Chomsky's influential distinction between competence and performance.<sup>8</sup> Then I will clarify the notion of *competent performance* at issue. This will put me in a position to characterize what I take to be the desiderata of a theory of competence.

### 2.1.1 Chomsky's Competence/ Performance Distinction

Chomsky (1965) aimed to provide a theory of linguistic competence—in particular, grammatical competence. But he was worried about trying to give an account of *all* linguistic performance. There are too many causal factors. You might, for example, fail to speak because you have laryngitis, or you might fail to understand what someone says to you because you aren't paying attention. Although these failures are failures to *do* something linguistic, the failures themselves aren't *linguistic* failures. We thus shouldn't expect a linguistic theory to explain one's failure on this sort of occasion. Nor, similarly, should we expect a theory of linguistic competence to explain the health of the throat or the paying of attention in cases of linguistic success.

This phenomenon has been noted elsewhere. For example, Ernest Sosa notes that

---

<sup>8</sup>Chomsky (1965). Here I will treat Chomsky as though he holds exactly the same view he did in 1965. That is an oversimplification, but harmless for my purposes.

in order for a performance to be competent—and so amenable to explanation by appeal to competences—it must not only be the case that the subject possesses the relevant competence (the competence must be *seated* in the subject), but she must also be in the right *shape* (awake, sober, etc.) and in the right sort of *situation* (good lighting conditions, etc.). All three “S”s are required in order for a performance to be competent—that is, in order for the performance to be the sort of thing that appeal to competence explains. Your bad driving is not a failure reflective of your competence to drive if it’s due to your being drunk, but rather is reflective of a failure of judgment. You’re in the wrong shape for exercising your competence. Your failing to hit a target in very windy conditions does not reflect upon your competence to shoot arrows, because you don’t have a competence to shoot arrows in that kind of situation.<sup>9</sup>

However, Chomsky (1965) did not just think that a theory of competence should make certain idealizations and restrict itself to aspects of the performances that seem to be properly linguistic; he thought that a theory of linguistic competence can and should be carried out in a way that wouldn’t tell you anything directly about the cognitive processing that underlies competent performances themselves. For example, consider a phrase structure grammar which takes

(1) The dog sleeps on the couch

and returns

(2) [<sub>S</sub> [<sub>NP</sub> [<sub>D</sub> The] [<sub>N</sub> dog]] [<sub>VP</sub> [<sub>V</sub> sleeps] [<sub>PP</sub> [<sub>P</sub> on ] [<sub>NP</sub> [<sub>D</sub> the] [<sub>N</sub> couch]]]]]

Whatever the theory says about how we get from (1) to (2), for Chomsky it is not supposed to be a specification of how, for example, hearers actually go from hearing (1) to understanding its grammatical structure. It is not supposed to tell you about the sorts of processes that underlie understanding what is said. It’s merely supposed to describe and predict the kind of syntactic structure that the (suitably idealized) subject does detect in understanding a sentence. Because this is the aim of the theory, Chomsky deems factors pertaining to how subjects actually competently perform

---

<sup>9</sup>See Sosa (2010, forthcoming).



“irrelevant”, for example “memory limitations, ... [and] errors (random and characteristic) in applying [one’s] knowledge of language in actual performance”.<sup>10</sup> This is so even though these factors, in contrast to those above, do sometimes seem to be responsible for *linguistic failures*, not just failures to do linguistic things.

There is a big debate about whether linguists should be engaged in Chomsky’s project, or rather in the project of explaining competent performance and so actual language processing. There is even debate about whether it is possible to provide the sort of abstract theory of competence that Chomsky wants to pursue.<sup>11</sup> I don’t want to get into these issues here. I just want to clarify what I am interested in. I am interested in the facts that explain competent performance—in this case speech production and understanding—and so I am interested in facts about memory limitations, characteristic errors, and so on that might help us to understand how we get things done so reliably.

In order to avoid getting bogged down in these issues, let’s distinguish two kinds of competences one might study. I’ll call the kinds of competences Chomsky is interested in “abstract competences”, and the kinds of competences I am interested in “performance competences”. Whereas abstract competences are not intended to explain how we get things done so reliably, performance competences are intended to do so.<sup>12</sup> Many cognitive scientists differ from Chomsky in their goals, and are interested in performance competences rather than abstract competences. What I have to say for the rest of the paper, where relevant to cognitive science, will be relevant to the science of competent performance and performance competence. When I just say “competence”, please assume it to mean *performance competence*.

---

<sup>10</sup>Chomsky (1965), p. 3

<sup>11</sup>See Scholz et al. (2011) for an overview.

<sup>12</sup>Clearly performance and abstract competences are not unrelated, for a theory of abstract competences would provide a useful starting point for investigating competent performance. (Chomsky agrees with this claim; see Chomsky (1965), p. 9.)

### 2.1.2 Which Are the Competent Performances?

We have so far narrowed our interest to theories of competence of the sort that explain competent performances. But which, exactly, are the competent performances? Minimally, they are the kinds of performances of the sort we began with—perceiving things as they are, effective linguistic communication, recognitional identification, and so on. We are interested in the natures of these perceptual, cognitive, linguistic, and agential achievements, and how we reliably accomplish them. So that we have a name for this kind of competent performance, I will call them “achievements”.

However, our intuitive sense of a competent performance does not just include achievements. There also seem to be *competent failures*—cases in which performance seems to be competent, but we fail to perceive things as they are, communicate, identify, and so on. Moreover, competent failures seem to be related to competent achievements in certain systematic ways. For example, we have visual illusions that seem to reflect features of the processing responsible for normal cases of perceiving things as they are. Consider the images of Figure 1.




Figure 2.1: The Müller-Lyer illusion and variations on it. From Howe & Purves (2005a).

(A) is the standard Müller-Lyer illusion. The two horizontal lines are the same length, even though the one on top looks to be longer than the one on the bottom. (B), (C), and (D) are all variations on the illusion. In each of these cases we misperceive the relative lengths of lines. This illusion is systematic, and we cannot learn to see the lines as being of the same length. This kind of illusion is often thought to provide clues as to how normal visual processing works. It seems that whatever cognitive mechanisms or processing explains our mistake in these cases is what also explains our seeing things as having the relative lengths that they actually have.

Competent failures, then, are standardly taken to fall within the domain of scientific theories of competence. Indeed, whether a scientific theory is interested in explaining competent failure is a good way of distinguishing theories of performance competence from theories of abstract competence. For example, the lab that created Figure 1 is interested in providing an account that explains features of our perceptual experience both in intuitively veridical cases and such illusory cases.<sup>13</sup> The philosopher would also do well, then, to include competent failures among the competent performances that her theory of competence explains.

### 2.1.3 Desiderata for a Philosophical Theory of Competence

I can now lay out some desiderata for a philosophical theory of competence. The main interest of a theory of performance competence (rather than abstract competence) is actual competent performance itself. We may divide this project into three tasks. First, the theory should, for any domain of competent performance, explain the sense in which the achievements and competent failures of that domain are both competent performances. Why do both of these kinds of performances count as *competent*?

Second, the theory should specify what competence possession is in a way that makes clear how appeal to a competence possessed by the subject can be illuminating of a particular case of competent performance. As noted above, central to the idea of competent performance is that competent performance is an exercise or manifestation of competence, and that features of the performance are *due to*, or *explained by* a competence possessed by the agent. This is also central to scientific practice, which aims to identify facts about the agent and her cognitive system that explain her performance on a particular occasion. Why is it that we better understand competent performance by appeal to certain stable properties of the agent?

---

<sup>13</sup>Howe & Purves (2005b), (2005a). See Purves et al. (2001) for an overview.

The first two desiderata require a theory of competence to provide a general framework for the relationship between achievements, competent failures, and the competences that explain them. The third desideratum I will impose on a theory of competence is that it be a useful framework for scientific investigation into competence possession and competent performance. At its best, a philosophical theory of competence will not only help us to better understand the person-level, everyday notions we are interested in (perceiving things as they are, inferring that  $p$ , understanding what is said, pouring coffee), but it will also help us to better understand how scientific characterizations of cognitive mechanisms and processes explain competent performance.

One might worry about this third desideratum. Does cognitive science need a philosophical theory of competence? If scientists are already engaged in the project of explaining competences and competent performance, what is there left for the philosophical theory to contribute to the project? We need a way of understanding how the rigorous (often formal) accounts of cognitive mechanisms provided by cognitive scientists relate to the larger question of explaining competent performance of human and other intelligent subjects. This is a philosophical project, although clearly not one that only philosophers are, or should be, engaged in. Getting clear on what competences and competent performances are will help us to interpret scientific theories and guide future research.

Sometimes what philosophical theory of competence is used as a framework for scientific investigation will not matter. I think there are several cases where conflicting theories of how scientific explanations explain competent performance will not have much impact on scientific investigation. In other cases, however, I think that having the right (or at least a better) account of competence will be very useful. In section 3 after I present my view of competences I will provide an example that shows how having a more efficient and fruitful competence framework for scientific investigation can help us to avoid scientific mistakes.

## 2.2 Comparison of Existing Approaches

Despite the centrality of competence to both scientific and philosophical investigation of intelligent performances generally, there has been little work directly on the nature of competences (or on capacities, abilities, powers, etc.). Nevertheless, a fairly clear pattern emerges of how achievements, competent failures, and competences are assumed to be related. We may call this approach the *success approach* for reasons that will become clear shortly. As we shall see, success approaches fail to explain the nature of achievements, and their relationship to competences.

Because of difficulties with the success approach, another approach is beginning to be explored, which I will call the *achievement approach*. While I think that the achievement approach corrects the major difficulty of the success approach, in its current form it is also severely flawed, failing entirely to explain competent failures. Instead, we need to develop a new theory that can adequately capture both kinds of competent performance.

### 2.2.1 The Success Approach

The success approach takes very seriously the first desideratum, that it must explain both achievements and failures. It aims to do so by fleshing out the following idea.<sup>14</sup> Perhaps there really is only one kind of competent performance, and it is *common* both to cases of achievement and cases of failure. Together with certain non-mental conditions, the competent performance constitutes an achievement, and together with others it constitutes a competent failure.<sup>15</sup>

So, for example, in the case of perception one might suppose the competent performance to be perceptually representing one's environment. Together with further non-mental factors one perceives one's environment as it is or one fails to do so. Or,

---

<sup>14</sup>It is worth noting here that philosophers often use the terms "ability", "capacity", and "competence" interchangeably. I do not think I am doing the views that I cite and discuss a disservice by considering them to employ a general approach to explaining competence, although some might deny that there is a reliability condition on competence possession. I will note this divergence where relevant.

<sup>15</sup>Here and throughout I will use the term "mental" somewhat broadly, to include any conditions that essentially involve subjects, persons, or agents, *as such*.

in the case of action, one might suppose the competent performance to be possessing an intention. When the intention causes the fulfillment of the intention in the right way, one intentionally does what one intends to do. Otherwise, one competently fails. (We'll come back to the question of what "in the right way" is supposed to mean.)

More precisely, the idea is this. We may specify competences in terms of a common exercise and a *success* condition that can be specified in terms of the exercise and non-mental conditions. Competences are competences to *succeed*, and achievements are to be explained in terms of the exercises of such competences, the obtaining of the success condition and further conditions relating the exercise of competence and the success condition.<sup>16</sup> Competent failures obtain when the competence is exercised but either the success condition fails to obtain or the exercise of competence and the success condition fail to be related in the right way.<sup>17</sup>

For example, on this approach perceptual competences might be competences to veridically perceptually represent. Such competences have as exercises cases of perceptually representing (which may or may not be veridical). When the fact that makes one's perceptual representation veridical appropriately causes one's having a perceptual representation, one perceives things as they are. While this approach is extremely popular—so popular it often goes unmentioned and undefended—I think it is wrong-headed. We can understand why by seeing how such accounts fail to deal with *deviant causal chains*. The way in which deviant causal chains are a systematic problem for the success approach reveals that it fails to accurately describe the relationship between competence and achievement.

---

<sup>16</sup>On some approaches, one only has a veridical representation if the external conditions constitutive of perception obtain. E.g. Burge (2010). See esp. pp. 379-381. This approach still aims to reduce achievements to common competent performances and non-mental conditions, and so I think will encounter similar difficulties. However, arguing for this is beyond the scope of this paper.

<sup>17</sup>For examples of those who employ the first approach, see, e.g., Bergmann (2008); Burge (2010, 2011); Greco (2010, 2012); Plantinga (1993); Schellenberg (2013); Sosa (2007). Only recently has this approach been articulated as a general strategy, by Ernest Sosa (forthcoming). See esp. p. 18.

### Deviant Causal Chains

According to the success approach, the subject achieves just in case she exercises her competence, succeeds, and her success is related “in the right way” to her competence. The problem is that whatever proponents of the success approach have proposed to be “the right way” for the exercise of competence to be related to success, there are counterexamples—it is possible to come up with a case where the conditions the theorist proposes are satisfied, but nevertheless the subject fails to achieve.

Veridical hallucinations, Gettier cases, and cases of deviantly fulfilled intentions are all cases of this sort. For example, one might suppose that *S* perceives an object (*o*) as it is (*F*) just in case *S* has a perceptual experience as of an *F* and *o*'s being *F* causes her to have a perceptual experience with that content. That this condition is insufficient is shown by the following case:

**Veridical Hallucination.** Sandra's new puppy has been giving her a headache. She reaches for pain killers but her friend Zoe, knowing how trying it is to have a new puppy, decides to play a trick on her, replacing her pain killers with pills that make one hallucinate black and white dogs. Sandra takes a pill, which causes her to have a perceptual experience as of a black and white dog on the floor before her. In fact, there is a black and white dog on the floor before her. (Her puppy is black and white).

This is a case of what some people call “veridical hallucination”—Sandra has a perceptual experience that “matches” a real-world scenario, but she hallucinates (purportedly) because the experience is not appropriately related to the fact that makes her experience veridical. This is so despite the fact that Sandra's puppy caused her to have the perceptual experience in question (via causing her to reach for the pain killers), and even is causally responsible for her having an “accurate” perceptual experience (via causing Zoe to decide to replace Sandra's pills with pills that cause hallucinations of dogs that look like Sandra's puppy).

The challenge for the success approach is to explain why the causal chain involved in Sandra's case is deviant—why the exercise of her competence is not related to success in the right way—but that does not overgeneralize and rule out cases when she actually does perceive a black and white dog as such. It may seem simple to block this

case, for example by requiring that light bouncing off of the object hit the retina and cause visual processing that causes the perceptual experience, but the problem will rear its ugly head again and again.<sup>18</sup>

There are also analogous cases for linguistic reference. Consider a variation on Kaplan's Carnap/Agnew case.<sup>19</sup> In the original case, the speaker points behind him to what he thinks is a picture of Rudolf Carnap and says, "That is a picture of one of the greatest philosophers of the Twentieth Century". In fact the picture of Carnap that usually hangs there has been swapped for one of Spiro Agnew, an American politician. In this sort of case Kaplan fails to secure what he intended as the referent of the demonstrative (e.g. the picture behind him which is a picture of Carnap). It seems that either the picture of Agnew is the referent or reference fails.

Consider now a variation of the case, in which the regular picture of Carnap has been replaced by a picture of Carnap dressed up as Spiro Agnew for a Halloween party. The costume is quite convincing, but the audience mistakenly thinks that *that* is what Carnap looks like. In this case, when Kaplan says, "That is a picture of one of the greatest philosophers of the Twentieth Century", Kaplan intends for the picture that is behind him which is a picture of Carnap to be the intended value, and the audience takes the picture (*qua* picture of Carnap) to be the value. However, that the audience "picks up on" the intended referent is merely lucky. The case is not the kind of achievement we normally take communication about demonstrated objects to be.

There is, of course, much more to say about this case. For example, the speaker may have a *de re* intention towards the original picture, not just Carnap, in which case the audience does not take the intended referent (the other portrait) to be the actual intended referent. Nevertheless, the worry remains. It is not enough that the speaker intend *o* to be the referent and that the audience take *o* to be the intended referent. There must be the "right" relation between the speaker's intention and the audience's taking *o* to be the intended referent. But whatever one supposes this relation to be, it seems possible to come up with a counterexample.

---

<sup>18</sup>See Noë (2003) for a brief overview.

<sup>19</sup>Kaplan (1989).



What, exactly, is the problem that deviant causal chains pose? Let me first say what the problem is *not*. It is not merely that the success approach fails to provide individually necessary and jointly sufficient conditions for achievements. There are very few successful metaphysical or conceptual analyses that do. The success approach would be in good standing if that were its main problem, for there are many interesting and fruitful *partial* metaphysical analyses, which merely aim to characterize some central aspects or structure of the facts in virtue of which the phenomenon of interest obtains. If the success approach is supposed to merely provide a partial metaphysical analysis, the failure to provide individually sufficient conditions is irrelevant.<sup>20</sup>

The reason why deviant causal chains pose a substantial problem for the success approach is that we can develop *procedures* for generating deviant causal chain cases. That there are such general procedures shows that *whatever* relation between the success and the common competent performance is proposed to be constitutive of achievement, it can obtain in a case that is intuitively not one of achievement. It is not merely difficult to find the right relation between the exercise of competence and the success—there is *no* relation that can do the job the success approach supposes is done.

I argue for this in detail for the case of knowledge in my “Competence to Know”.<sup>21</sup> The case of knowledge is instructive because there have been so many attempts to solve the Gettier problem and so the procedure can be tested against several existing well-developed theories, but the thought behind the procedure and the features of indirect virtue epistemology that it depends on are fully general.

To get a sense for how this goes, consider a version of Chisholm’s sheep Gettier case:

**ROCK:** Annette is taking a walk through the countryside. Looking at what seems to her to be a sheep in the field, she forms the belief that there is a

---

<sup>20</sup> Sosa (forthcoming) proposes that proponents of the success approach (my terminology) treat their explanations of achievements as partial metaphysical explanations, supposing that the exercise of competence must be related to the success “in the right way”, although exactly what that is remains inarticulable.

<sup>21</sup> Miracchi (2014b). I present a procedure for generating counterexamples to what I call indirect virtue epistemology, which is an application of the success approach to the case of knowledge.

sheep in the field on the basis of her perceptual experience. In fact it is a sheepdog, but there is a sheep standing behind a rock, out of view, that the dog is keeping track of.

ROCK is a Gettier case, because intuitively Annette has a justified belief and she believes truly, but nevertheless she fails to know. To put it in competence terms, we may say that Annette exercises a competence to believe truly, she does believe truly (succeeds), but it her success isn't related to her exercise in the right way, i.e. in the way constitutive of knowledge.

Now, suppose someone (in this case Greco (2012)) comes along and says that Annette fails to know because her competence isn't involved in a way that would "regularly serve [the subject's] relevant informational needs".<sup>22</sup> That is, the particular way in which the competence is involved is not one that would reliably get Annette true beliefs. We can then emend the case in the way just suggested to ensure that Annette's competence is involved in a way that regularly serves her relevant informational needs—i.e., in a way that reliably gets Annette true beliefs. We do this by establishing a connection between what is intuitively the "bad" and "good" luck of ROCK, leaving intact whatever facts are intuitively responsible for competence possession. For example, we can use ROCK to generate the following case:

**ROCK AND HARD PLACE:** Annette is taking a walk through the countryside. Looking at what seems to her to be a sheep in the field, she forms the belief that there is a sheep in the field on the basis of her perceptual experience. In fact it is a sheepdog, but there is a sheep standing behind a rock, out of view, that the dog is keeping track of. Unbeknownst to Annette, she is in hard-working sheepdog country, in which it is very rare for a sheepdog to be in a field unless it is keeping close watch on its sheep.

ROCK AND HARD PLACE is a case where Annette's competence reliably serves the purpose of forming true beliefs, because it is highly likely that whenever she believes that there is a sheep in the field on the basis of seeing something which looks to her to be a sheep in the field (but is in fact a dog) her belief is true (because the dog is watching a nearby sheep). Thus Rock and Hard Place satisfies Greco's conditions, but nevertheless Annette does not know.<sup>23</sup>

---

<sup>22</sup>Greco (2012), p. 17.

<sup>23</sup>Importantly, this case is not a fake barn case (Goldman (1976)). Annette's perceptual faculties are

Note that all this procedure relies upon is the idea that whatever general relation between the exercise of competence and the success is supposed to be “the right way”, it can be imposed on a deviant causal chain case by establishing a certain kind of relation between the “bad” luck (e.g. a dog looking like a sheep) and the “good” luck (a sheep being behind a tree) of the deviant causal chain case. Thus the procedure doesn’t depend on anything particular to the epistemic domain, but rather on the central features of the success approach.

We can understand why such procedures exist for generating deviant causal chain cases if we reject the starting assumption of the success approach, namely that achievements can be explained in terms of non-factive exercises of competences to succeed and non-mental further conditions. Instead, achievements are competent performances in their own right: they are exercises of competences to *achieve*.

So, for example, perceiving things as they are is itself an exercise of a competence to perceive, and is not a competent performance merely because it is constituted by an exercise of a competence to veridically visually represent together with further conditions. Referring to an object with a demonstrative is an exercise of a competence to refer with a demonstrative, and is not a competent performance merely because it is constituted by an exercise of a competence to have one’s audience take the intended object to be the intended object together with further conditions.

What makes a causal chain “deviant” on this way of looking at the matter is that it deviates from what the competence is a competence to do. That is, it deviates from cases of perceiving things as they are, referring with a demonstrative, etc. If *that’s* right, then we can start with a case that intuitively has a deviant causal chain, and no further non-mental changes or conditions that we add to it will turn it into a case of achievement. Once we’ve deviated from the case of achievement (e.g. perception or reference), we’re sunk. However, that is not what the success approach predicts. The success approach claims that it is certain general, competence-independent relations

---

reliable at discriminating sheep from non-sheep in hard-working sheepdog country. It is just that, when she is in circumstances where a dog in a field looks like a sheep to her, it is highly likely that her perceptual belief that there is a sheep in the field will be true.

between exercises of competence and success that makes an achievement obtain. We can thus take a deviant causal chain case, add in whatever conditions the success approach claims are constitutive of achievement, and generate a counterexample to the view.<sup>24</sup> This is precisely the procedure we followed in generating ROCK AND HARD PLACE from ROCK.

Although I do not have the space to argue that the approach generalizes in more depth here, it should be clear that minimally the proponents of the success approach face a compelling challenge. They must either explain why in these other domains the procedure will not generalize, or they must give good reason for thinking that this isn't a serious problem. For now, I will suppose that they will not be able to satisfactorily answer this challenge and will move on to the other main approach to competences, the achievement approach.

### 2.2.2 The Achievement Approach

Recently, there has been some backlash against the success approach, largely in response to doubts about its ability to respond effectively to deviant causal chain cases.<sup>25</sup> Instead of thinking that competences are competences to succeed that have neutral exercises, some have supposed that competences are competences to achieve that have achievements as their exercises. While I think that this general approach is correct, in existing forms it is seriously deficient. This is because such accounts either claim that *all* exercises of the competences that explain achievements must be achievements, or they ignore the task of showing how competent failures might also be exercises of the same competences.

That the need to explain competent failures is a pressing philosophical concern has not been adequately appreciated by those who reject the success approach. Alan

---

<sup>24</sup>To be clear, the objection is *not* that the success approach invokes causal facts to explain achievements. (As we shall see, my account also does this.) The objection is that the causal facts that explain achievements must also explain the sense in which achievements are performances by subjects—not just products (or causes) of their competence, but also performances themselves. The success approach does not do this by design, and so there will always be cases that satisfy such accounts but fail to be achievements.

<sup>25</sup>See, e.g., Williamson (2000).

Millar and John McDowell are two proponents of the achievement approach who discuss competences (abilities/capacities) in some depth. Millar flatly claims without argument that exercises of competences to achieve must *always* issue in achievement:

If I had judged falsely that the plants in the plot were azaleas I would not have exercised the recognitional ability in question. The general point here is that the notion of the exercise of a recognitional ability is a success notion. Success in relation to finding out whether or not *p* is coming to know whether or not *p*.

—Millar (2008), pp. 3-4

Here Millar assumes that any exercise of an ability to know on the basis of perceptual experience that an azalea is before one just is a case of coming to know on the basis of perceptual experience that an azalea is before one. This idea precludes explanations of competent failures (e.g. visual illusions as of azaleas, and false justified beliefs formed on that basis) in terms of the same competences that explain cases of knowledge. But both intuitively and scientifically, it seems that the same competences *do* explain these two kinds of case.

McDowell concedes that perceptual capacities may be fallible, but is unmoved to explain how.<sup>26</sup> He claims that the possibility of the fallibility of our perceptual capacities only poses a problem to the view that its exercises are cases of perceiving things as they are if it shows that the idea of a factive mental state of seeing things as they are is unintelligible.<sup>27</sup> However, this is just not true. Many opponents of disjunctivism about perception (and McDowell's view in particular) explicitly claim that their position is to be preferred because it provides explanations of illusions and hallucinations, on which McDowell and other disjunctivists are notably silent.<sup>28</sup> They are right to do so. A theory of perceiving things as they are (or any other achievement) should also explain its relation to corresponding competent failures.

---

<sup>26</sup>E.g. McDowell (1994), p. 112; McDowell (2010), p. 3.

<sup>27</sup>"It would matter if it showed that the very idea of openness to facts is unintelligible, and it does not show that.... The aim here is not to answer skeptical questions, but to begin to see how it might be intellectually respectable to ignore them..." (McDowell (1994), p. 113).

<sup>28</sup>E.g. Burge (2005); Schellenberg (2010).

## 2.3 Fallible Competences, Factive Exercises

### 2.3.1 A Middle Way?

Note that we can organize the success approach and the achievement approach in the following table:

	Fallible Competence	Infallible Competence
Non-Factive Exercise	Success Approach	
Factive Exercise		Achievement Approach

The success approach holds that competences are fallible and that they have non-factive exercises (e.g. competence to veridically perceptually represent, perceptually represent). The achievement approach holds that competences are infallible and their exercises are factive (e.g. competence to perceive things as they are, perceiving things as they are)—or at the very least is unmoved to explain how it is possible for competences with factive exercises to be fallible. Underlying this division of positions on competences is the (often implicit) agreement that *factivity* of the exercise and *infallibility* of the competence come as a package—if an account supposes that achievements just are exercises of competence, the competence may never issue in exercises that are failures.<sup>29</sup>

But it is at least a *logical* possibility that competences should both be fallible and have factive exercises. Is this a live possibility? Can we understand how competences could be competences to achieve, have some exercises that are constitutively achievements, and also to have some exercises that are constitutively failures?<sup>30</sup> The best

---

<sup>29</sup>Note that I am here using the term “factive” loosely to include case of, e.g., demonstrative reference, even if there is no propositional content of the state that must be true in order for the state to obtain.

<sup>30</sup>The view is evident in, among others, the work of Sosa (2007), Burge (2010), Burge (2011), Millar (2008). McDowell is often read as requiring that perceptual capacities be infallible because he holds that perceiving that *p* provides the subject with indefeasible warrant to believe that *p*. This is a very natural picture of justification, and is a central piece of Goldman’s (1979) reliabilist theory of justification. But in his (2010) he makes clear that that is not the view. He must hold that degree of justification does not derive from the competence responsible for the experience but the experience itself, and proper basing of that experience. I prefer Goldman’s approach, but of course the other approach is standard as well,

existing argument against this view is a modal argument, made by Tyler Burge (2011).

### 2.3.2 The Modal Argument Against Fallible Competences with Factive Exercises

Burge makes the argument by using an analogy with a capacity to make free throws:

To realize the function of the capacity to shoot free throws is to actually make them. But any given successful shot could have been unsuccessful if circumstances had been different. The very same shooting that, in fact, was successful could have failed to be successful—for example, if immediately after the individual shot the ball, a sudden wind had blown it off course. That is a metaphysical possibility. ... Particular perceptions seem to be subject to the same possible mismatches that make the general capacity fallible. Nothing about the perceptual ability or its exercise on the particular occasion guarantees veridicality.

—Burge (2011), p. 54

We can reconstruct the argument as follows:

#### BURGE'S ARGUMENT

1. In a case where I exercise my competence and in doing so achieve, it is true that I could have exercised my competence and in doing so have failed.
2. By (1), a particular exercise of competence in a case of achievement could have occurred in a case of failure.
3. Therefore, a particular exercise of competence in a case of achievement does not entail that an achievement obtains (but is instead non-factive).

I don't think this argument is a very good one. To see this, consider an analogy. It's a toy example, but you get the point:

#### JADE ARGUMENT

1. In a case where I give you jade and in doing so give you jadeite, it is true that I could have given you jade and in doing so given you nephrite.<sup>31</sup>

---

defended by Earl Conee and Richard Feldman (1985).

<sup>31</sup>Jade is a disjunctive kind. There are two kinds of rocks that are jade, jadeite and nephrite.

2. By (1), a particular giving of jade and in doing so giving jadeite could have occurred in a case of giving nephrite.
3. Therefore, the giving of jade in the case of giving jadeite does not entail giving jadeite (but is instead not object-involving).

(Please treat “giving of” extensionally.)

Now *this* argument is clearly not any good, because the giving of jade in a particular case does entail giving an object that is jade (in this case the particular piece of jadeite). There is no giving of jade in the case under consideration that does not involve the jadeite. The fact that one could have given jade and not given jadeite doesn't mean that there is a non-object-involving giving in either the jadeite or nephrite cases.

Here is the problem with Burge's argument. Burge assumes that facts about counterfactual properties of an actual scenario reveal what is essential about the things or events involved in that scenario. But when we are working with disjunctive properties, that is not true. A disjunction obtains if either of its disjuncts obtains, and so neither disjunct is necessary for the disjunction. But that at least one of the disjuncts obtains *is* necessary for the disjunction to obtain. We cannot infer that giving a particular kind of jade is not necessary for giving jade at all from the fact that one could give jade of one kind or another.

If, as I am suggesting, *exercise of competence* is a disjunctive kind, then we cannot infer from the fact that neither achievement nor competent failure is necessary for an exercise of competence to obtain that it is not necessary for at least one of them to obtain. We may not infer that competences may be exercised in ways that do not entail whether whether an achievement or competent failure obtains. Thus Burge's argument is invalid.

Burge concedes that his argument is really only an argument from inference to the best explanation:

I think that the inference from the fallibility of perceptual capacities to the fallibility of their exercises is a natural one. But I do not think that the inference rests on an entailment. Rather, I reflect on the contingencies



in causal connections to objects that successful instances of perception depend on, and connect those contingencies to the contingencies that make the perceptual capacities fallible. I judge it possible for a successful perception to have been unsuccessful if those contingencies had not been in place.

—Burge (2011), p. 53

An argument to the best explanation can be adequately dealt with by providing a competing explanation. If we can make clear how competences could be fallible and nevertheless have some factive exercises, then we would have a better explanation of competent performance, because we both explain fallibility and avoid deviant causal chains.

### 2.3.3 The Dual Exercise Account

We need an account of the nature of competences and their exercises that makes it plausible that competences are competences to achieve, and that their characteristic exercises just are achievements, but that they can have *degenerate exercises* too—exercises that are constitutively failures. In that way, we can both capture the idea that achievements are competent performances in their own right and the idea that the same competence explains both achievements and their intuitively corresponding competent failures.

In order to do this, we do not need to provide necessary *and* sufficient conditions for competence possession and exercise. It will suffice to impose certain merely necessary conditions—a partial metaphysical explanation, as discussed above. The tools we need to specify such conditions are actually quite mundane, and for the most part already in use.

The success approach to competences imposes a *reliability* condition on competence possession that relates the exercise of a competence and its success condition. For example, one might suppose that a subject has a competence to veridically perceptually represent only if there is a high enough probability that she will veridically perceptually represent given that she perceptually represents. More generally, where

$S$  is the obtaining of the success condition and  $E$  is the exercise of the competence to succeed:

SUCCESS APPROACH RELIABILITY CONDITION.

$$Pr(S|E) \geq n, \text{ for some sufficiently high } n \in (0, 1].^{32,33}$$

This condition states that when the subject exercises her competence, it is likely that she will succeed. If the success approach could overcome the difficulties discussed in the last section, this reliability condition would help to explain why exercises of competence are *competent* performances—when the competence is exercised, the subject has a high probability of succeeding.

Instead, we can impose a different kind of reliability condition, one that will make it likely that whenever the subject exercises her competence she will achieve. We *could* specify a reliability condition just by replacing the “ $S$ ” above for an “ $A$ ” (representing a case of achievement), to get the following condition:

UNHELPFUL RELIABILITY CONDITION #1.

$$Pr(A|E) \geq n, \text{ for some sufficiently high } n \in (0, 1].$$

This condition, however, is not very helpful. So far, it is equivalent to the requirement that there be a high probability that  $A$  obtains given that  $A$  or  $F$  obtains, where “ $F$ ” represents an intuitively corresponding case of competent failure:

UNHELPFUL RELIABILITY CONDITION #2.

$$Pr(A|(A \vee F)) \geq n, \text{ for some sufficiently high } n \in (0, 1].$$

---

<sup>32</sup>For ease of exposition, I am representing the threshold as a real number  $n \in (0, 1]$ —that is,  $n$  is a number between 0 and 1, including 1 but not 0. However, I need not make any commitments about how close to 1  $n$  must be, whether the threshold for the epistemic domain is the same as for other domains (e.g. baseball), or whether the threshold is determinate. (I think it is probably indeterminate.) I also need not make any commitments about whether the threshold is context-sensitive.

<sup>33</sup>While I need not commit to a particular view of probability here, I am supposing that appeal to objective conditional probabilities is less contentious than appeal to objective unconditional probabilities (see Hájek (2007) for defense of this claim), and I am supposing that some kind of non-frequentist realist account of them is in the offing. I am also supposing that such probabilities are true at particular times. Changes in, e.g., causal regularities over time might result in either the acquisition or loss of competences to know by changing whether or not the proficiency condition is met.

Unlike the success approach reliability condition, these unhelpful reliability conditions do not explain why achievements and competent failures are both performances of the same competence, and thus why it matters to whether or not they are *competent performances* that when either the subject achieves or competently fails there is a high probability that she achieves.

To see the problem, note that the fact that my dog Owen has a high probability of barking given that he either barks or meows is not informative as to why he is competent at barking. Meowing isn't related to barking in the way that would make this reliability condition illuminating. We thus face an explanatory burden: why should it matter to understanding why either achievements or their corresponding failures are competent performances that it is highly likely that if either of these cases obtain the subject achieves?

A view that only invoked the unhelpful reliability conditions would fail to satisfy desideratum #1 on a theory of competences—it would fail to explain the sense in which achievements and competent failures are both competent performances. How can we do better? Instead of discharging this burden by supposing that that exercises of competence are already specifiable *qua* performances independently of their being involved in cases of achievement or competent failure (believings, representings, intentions, etc.) as the success account does, we may look to the facts *in virtue of which* subjects exercise their competences in order to specify our reliability condition. That is, we may return to one of the ideas we began with, namely that the *nature* of a performance might be determined by the facts that make it a *competent performance*. Recall the examples of representation and perceptual representation given above (p. 2). There the idea was that what made it the case that a subject represented at all made it the case that she competently represented. Although in that case competent representation is characterized in terms of a success condition (veridicality), this is not the only way in which the idea may be developed.

By holding that the probability facts constitutive of competences are also constitutive of the nature of achievements directly, we can eliminate the need to explain

achievements in terms of other mental states and activities that are non-factive. Instead, our reliability condition relates certain facts that are constitutive of achievements to the operations of subpersonal cognitive (and perhaps bodily) facts that we may suppose to be the *basis* of the competence, which is seated in the subject.<sup>34</sup> In particular, we may begin by defining the *manifestation conditions* of a competence  $C_A$  to  $A$  as follows:

MANIFESTATION CONDITIONS: The manifestation conditions of  $C_A$  are whatever operations of subpersonal cognitive (and perhaps bodily) mechanisms and external conditions together (against a background of possession of  $C_A$ ) constitute a particular case of  $A$ -ing.

A competence  $C_A$  to  $A$  is MANIFESTED just in case its manifestation conditions obtain.

Of the manifestation conditions, we may isolate just the subpersonal cognitive (and perhaps bodily) mechanisms whose operations in a particular case partially constitute the subject's  $A$ -ing:

BASIS CONDITION: The basis of  $C_A$  is fully constituted by the subpersonal cognitive mechanisms of the subject  $S$  whose operations partially constitute  $S$ 's  $A$ -ing.

Although these two conditions are not very informative, they allow us to define the kind of reliability condition we are looking for:

PROFICIENCY CONDITION: The proficiency condition of  $C_A$  requires that the objective probability of the manifestation conditions obtaining conditional on the basis of the competence being operative be sufficiently high. I.e.,  $Pr(M|O_B) \geq n$ , for some sufficiently high  $n \in (0, 1]$ .

---

<sup>34</sup>Depending about which competence is in question, sometimes it may be plausible that certain bodily movements, etc., are part of the operation of the basis of the competence, not just the manifestation conditions. For perception, thought, and so on, it is plausible (though not uncontested) that the basis resides solely "in the head". For cases of intentional action and certain aspects of linguistic competent performance, the basis will likely be broader, perhaps involving various motor movements. I have in mind here competences to make free throws, to hit a bull's eye, etc. I leave open here exactly where the line between the operation of the basis and the rest of the manifestation conditions should be drawn.

Because the proficiency condition allows that the probability of the manifestation conditions obtaining conditional on the basis being operative may be *less* than 1, it allows that the bases of competences responsible for achievements may be operative even if the full manifestation conditions fail to obtain. We may now suppose that a competence is *degenerately exercised* just in case its basis is operative but not all of its manifestation conditions obtain. These are cases of competent failure.

We now arrive at a theory of competences on which they have exercises that are achievements, but are nevertheless fallible. Because the reliability condition characteristic of competences relates the facts that in a particular case are constitutive of achievement to their cognitive (proper) part, and does not appeal to any other competent performances, we can understand how competences can have two kinds of exercise: manifestations, which just are achievements, and degenerate exercises, which just are competent failures. Here is the full account:

#### DUAL EXERCISE ACCOUNT OF COMPETENCES

1. For any achievement  $A$ , any competence  $C_A$  to  $A$  has a *basis* and *manifestation conditions* which satisfy the *proficiency condition* just defined.
2.  $C_A$  is MANIFESTED just in case its manifestation conditions obtain. A manifestation of competence  $C_A$  is a case of  $A$ -ing, an ACHIEVEMENT.
3.  $C_A$  is DEGENERATELY EXERCISED just in case its basis is operative, but not all of the manifestation conditions obtain. A degenerate exercise of  $C_A$  is a COMPETENT FAILURE to  $A$ .
4.  $C_A$  is EXERCISED just in case it is either manifested or degenerately exercised.

The first thing to note is that the conditions just stated entail our unhelpful reliability conditions. When the basis of a competence is operative, either the rest of the manifestation conditions obtain or they fail to obtain. That is, either the subject achieves or she competently fails. Since the proficiency condition requires that it be highly likely

that whenever the basis of a competence is operative its manifestation conditions obtain, it also ensures that whenever the subject exercises her competence it is highly likely that she achieves.

However, unlike the unhelpful reliability conditions, there is a straightforward explanation of why the obtaining of the proficiency condition matters as to whether achievements and competent failures are competent performances. Only if the proficiency condition obtains do manifestation conditions constitute an achievement at all, for this is a necessary condition on possession of the competence that was manifested. (Hence the caveat about background possession of the competence in the statement of the manifestation conditions.) Thus only if an achievement is competent in the sense of being likely to be an achievement (*qua* exercise of competence) is it an achievement at all.

Similarly, only if the proficiency condition obtains does the operation of the basis in conditions that deviate from manifestation conditions constitute a degenerate exercise of competence. That is, a competent failure is the kind of performance it is only because it deviates from a case of achievement. Thus only if a competent failure is likely to be an achievement (*qua* exercise of competence) is it a competent failure at all.

To sum up so far: the dual exercise account explains how the competences which explain achievements and competent failures can be fallible and yet have (some) factive exercises. In doing so, it also satisfies the first desideratum on a theory of competence, explaining how achievements and competent failures are both competent performances. How does it fare with respect to the other two desiderata?

*Desideratum #2: Why Is Competence Possession Illuminating of Both Achievements and Competent Failures?*

There are two different ways in which the success approach answers this question, one of which the dual exercise account rejects and the other of which it accepts. It rejects the idea that competences may be specified independently of the achievements

they purportedly explain. On the success approach, the competences which explain achievements are competences to succeed. They are fully specified in terms of a non-factive exercise and a success condition. It is a *further* fact, independent of the nature of the competence in question, what it is for a success to be due to competence in the way constitutive of achievement. One way in which competences explain achievements on this approach, then, is that they are part of a reductive analysis of them.

According to the dual exercise account, in contrast, what it is to *A* is to be a manifestation of a competence to *A*. *A*-ing and possession of a competence to *A* are metaphysically *co-dependent*, and the facts that determine possession of a competence to *A* determine what it is to *A*. So the dual exercise account does not aim to explain achievement in terms of independently specifiable competences, but rather aims to explain achievements and the competences they manifest in terms of more basic facts.

However, it shares with the success approach the idea that the nature and distinctive features of competent performances is determined by the competences they manifest. Thus by investigating features of competence possession and the facts in virtue of which subjects possess competences, we can better understand competent performances.

Indeed, on the dual exercise account, the operations of bases of competences are partially constitutive of both kinds of competent performances. Thus investigation into these bases as well as the regularities constitutive of competence possession will help us to understand both achievements and competent failures, as well as what they have in common.

*Desideratum #3: Does the Dual Exercise Account Provide a Useful Framework for Scientific Explanations of Competent Performance?*

On the dual exercise account, competent performance can be investigated by studying the basis of the competence responsible for the performance and the regularities relating aspects of this basis to features of the subject's environment. Such investigation

can (and should) be pursued scientifically. Cognitive scientists can help us to understand the nature of the bases of competences and the facts constitutive of competence possession by investigating the operations of cognitive mechanisms and their relations to the environment. This way of conceiving of how scientific investigation illuminates competent performance is a much more flexible, but still informative, framework for scientific investigation into competent performance than the success approach.

First, it allows for a distinction between conceiving of the brain as an information-using system without conceiving of the *mind*, or mental states, as such. The cognitive mechanisms that are the object of scientific study are the bases of competences; the manifestations of such competences are mental states and activities. This approach, while *compatible* with the popular view that mental states are representational states and that they are identical with representational states of cognitive systems, does not require it.<sup>35</sup> It allows there to be a more complex relation between the cognitive mechanisms that are the object of scientific study and person-level competent performance, while still maintaining a close connection between them.<sup>36</sup>

This is crucial if a philosophical theory is to serve the purpose of providing a framework for cognitive science. While the field has made important strides in the last few decades, it is still young, and fundamental questions remain hotly debated. A general philosophical framework for the study of competent performance should provide a ground for debating these questions, and should not force a particular kind of answer where there is not sufficient philosophical or empirical reason to do so. I believe there is *not* sufficient reason to suppose that mental states are representational states and that they are identical with representational states of cognitive systems: we may be able to explain how the operations of cognitive mechanisms and environmental facts underlie competent person-level performance without making this commitment.

---

<sup>35</sup>More precisely, it is compatible with certain versions of the view, on which, e.g., what it is for a subject to perceive is to veridically visually represent. (This is Burge (2010)'s view.) On this view, there is no gap between the success condition and the achievement condition, and so the problem of deviant causal chains as posed above does not apply. Nevertheless, I think that the considerations presented above do put pressure on this approach. Arguing for this, however, would take us too far astray.

<sup>36</sup>This idea of how computational explanations explain competent performance has also been recently defended by Egan (2013).



Dennett provides a nice example of this in “A Cure for the Common Code?”:

In a recent conversation with the designer of a chess-playing program I heard the following criticism of a rival program: “It thinks it should get its queen out early.” This ascribes a propositional attitude to the program in a very useful and predictive way, for the designer went on to say, one can usually count on chasing that queen around the board. But for all the many levels of explicit representation to be found in the program, nowhere is anything roughly synonymous with “I should get my queen out early” explicitly tokened. The level of analysis to which the designer’s remark belongs describes features of the program that are, in an entirely innocent way, emergent properties of the computational processes that have “engineering reality”. I see no reason to believe that the relation between belief-talk and psychological process talk will be any more direct.

—Dennett (1978), p. 107<sup>37</sup>

While Dennett here uses the competent performance as reason to attribute a thought *about* the competent performance, we may take the example to merely provide a lesson about the relationship between cognitive processing and competent performance. The crucial thing is that Dennett is supposing (as I think is quite plausible) that we might have a computational explanation of why the chess program reliably (competently) gets its queen out early that does not posit any state of the representational system that represents that it should get its queen out early.

Rather, the explanation of why the robot reliably gets its queen out early must appeal to both features of the cognitive processing and its relation to features of the environment. If perceptual, intentional, and other person-level performances are competent performances, they may likewise be explained by subpersonal computational processing even if that processing includes no symbols with the content that the person-level performances intuitively have.

While adopting the dual exercise account is not the only way to accommodate this possibility, doing so provides a clear and straightforward account of the relationship between states and processes of cognitive systems and the competent performances they explain.<sup>38</sup> The operations of bases of competences do not themselves constitute

---

<sup>37</sup>Fodor (1987) discusses this passage on pp. 21-3.

<sup>38</sup>It also does not undermine realism about these performances as some, e.g., Dennett’s way of accommodating this example does.

competent performances; only together with environmental factors (the full manifestation conditions or deviations from these conditions) do competent performances obtain. Thus there is no pressure to identify the intentional states of persons with states of cognitive systems, much less their contents with the contents of cognitive systems. Instead, features of cognitive processing together with features of the environment explain why a subject perceives things to be thus and so, or thinks that *p*.

This is not merely a hypothetical upshot; there are already good computational explanations of competent performance that do not posit states of the cognitive system with contents that appropriately correspond to the competent performance explained. Here is an example from current scientific work on human visual performance. Recall the Müller Lyer illusion and related illusions (Fig. 1) from the beginning of the paper. Dale Purves and his lab have proposed one of the best current scientific explanations of why we perceive the relevant lines as being of different lengths. This theory does *not* involve a computational process resulting in the attribution of different lengths (or of relative difference in length) to the lines. Rather, it supposes that cognitive systems compute highly abstract statistical information about retinal stimuli and features of our phenomenal awareness correspond to these statistical properties.<sup>39</sup> This approach is highly predictive of the phenomenal character of experience, including many other kinds of visual illusions.<sup>40</sup>

However, Purves' lab has interpreted their theory as having the result that we don't really perceive or misperceive the relative lengths, as well as a host of other features of objects—indeed, they conclude that the whole idea that we perceive objects as they are is mistaken:

Although we routinely attribute visual perceptual qualities to objects and environmental conditions, our experiences of lightness, brightness, color, form, and motion are likewise subjective qualities that simply promote useful behavior. Accordingly, it would be best to describe visual perceptions in terms similar to those used to describe pain, for which the concept of representation makes no sense. Visual perceptions, like the perception of pain, do not stand for the properties of objects in the physical world,

---

<sup>39</sup>Howe & Purves (2005b,a); Purves et al. (2011, 2001); Purves & Lotto (2002).

<sup>40</sup>For an overview see Purves & Lotto (2002).

although the world, of course, generates the relevant stimuli.

—Purves et al. (2011), p. 6

It is a radical claim that we do not see things as they are (and competently so) in visual perception—that visual perception is mere sensation, as Purves et al. assume pain to be. But it is not difficult to see why they draw this radical conclusion. If one supposes (as is often supposed) that visual processing produces a representational state that is the perceptual experience and that how things seem to one is determined by the content of the representational state—then whatever representational states posited by the scientific theory co-vary most closely with changes in phenomenal character will be the perceptual experience. If this were so, Purves' results would mean that we have perceptual experiences as of objects having certain statistical properties, which we clearly do not. Thus the alternative option—that the properties represented by visual systems that determine the phenomenal character of experience do not determine anything more than sensational qualities, may be preferred.

At this point, what should *not* happen is the rejection of a highly predictive scientific theory because it, together with certain *philosophical* assumptions about the relationship between representations of the visual system and visual experience, predicts implausible things about the nature of visual experience. What should happen is a re-assessment of the philosophical assumptions about the relationship between the rigorous scientific explanation and perceiving things as they are.

Because the dual exercise account does not suppose that there is a perceptual state that is fully determined by features “inside the head” together with certain causal regularities, it avoids this problem. From a scientific point of view, what is required is not looking for a different interpretation of the content of computational symbols in the visual system, but instead an explanation of how symbols with that content play the appropriate role in the visual system, together with features of the environment, to constitute our visually perceiving objects in accordance with the features that they have.

If Purves et al. are right, doing anything else would lead us on a *scientific* wild goose chase, as they are well aware:

The distinction between visual perception conceived in terms of the awareness of behaviorally useful qualities vs. conceiving perception in terms of representations of the physical world would be a philosophical point only, were it not for the associated neurobiological implications. If vision does not represent the properties of objects and conditions in the world, then neither do its underlying anatomical and physiological mechanisms, which must, therefore, be thought of, examined, and tested in different terms.

—Purves et al. (2011), p. 6

This is just one among many examples of how the dual exercise account provides a usefully flexible framework for understanding the relationship between the cognitive mechanisms that underlie competent performance and the features of competent performances themselves, one that allows us to avoid unnecessarily constraining scientific explanation or improperly directing it.

## 2.4 Conclusion

I have argued for the importance, both to philosophy and to cognitive science, of developing an accurate and useful theory of competences. I have argued against current existing approaches and developed and defended a theory which is both philosophically and scientifically more plausible. Perceiving, learning, walking, thinking, speaking, driving, understanding what is said, pouring coffee, inferring, and so on, really are kinds of competent performance, and need not be understood in terms of a non-factive, non-situation-involving, performances. The dual exercise account provides a framework for investigating these achievements and their corresponding failures directly. Of course there is much more to explain, both about competent performance in general and about its application to particular domains. Here I discuss one of the remaining questions.

I have provided a theory of competence by imposing necessary conditions on their possession and exercise, but I have said little about what competences are. Are they kinds of dispositions?<sup>41</sup> Are they *sui generis*? What kind of property is competence

---

<sup>41</sup>This is the view held by Sosa (2010).

possession? Here is one reason to think that competences are not merely a species of disposition, but instead are *sui generis*.

Dispositions are often thought to have associated *trigger-manifestation* conditionals.<sup>42</sup> For example, a match has a disposition to light when struck. “If the match were struck, it would light” is the corresponding trigger-manifestation conditional. The antecedent is a trigger, or cause, of the manifestation. The conditional is (more or less) true of the match because it has the disposition it does. Of course there are cases in which the associated conditional is false and the match still possesses the disposition (e.g. in a situation with no oxygen), but there does seem to be an important connection between possession of a disposition and the accuracy of such a conditional.

Associated with competences, on the other hand, is another kind of conditional, what we may call an *exercise-achievement* conditional, or perhaps an *aim-achievement* conditional. For example, if a subject has a competence to refer with a demonstrative, then “if she were to exercise her competence to refer with a demonstrative, then she would so refer” is an associated conditional. So is, “if she were to aim to refer with a demonstrative, then she would so refer”. To give another example, suppose that someone has a competence to make free throws. If this is so, then “if she were to shoot, she would make the free throw” is an associated conditional, as well as “if she were to try to make the shot, she would make it”. There are of course exceptions in this sort of case as well, but the associated conditionals do seem to capture something important about competence possession.

One might think that the conditionals associated with competences are really just a species of the trigger-manifestation conditionals associated with dispositions. This would give us reason to think that competences just are dispositions, and the way in which they explain their manifestations is the way in which dispositions explain their manifestations.<sup>43</sup>

However, the preceding discussion gives us good reason to think that the relationship that exercises bear to achievements is *not* causal. If we assume the relationship

---

<sup>42</sup>See Choi & Fara (2012) for an overview.

<sup>43</sup>Ernest Sosa (2010, forthcoming) makes precisely this proposal.

to be causal, then deviant causal chains will arise. Rather, *exercise of competence* is a disjunctive kind, and achievements are one of the disjuncts. The conditionals associated with competences are true (more or less), not because exercises tend to appropriately cause achievements, but because the proficiency condition makes it probable that an exercise of competence *is* an achievement. This suggests that the way in which competences explain their manifestations cannot be assimilated to the way in which dispositions do.

Instead, if the dual exercise account is right, what explains these conditionals is the relationship between subpersonal cognitive systems and environmental facts. These of course involve causal facts, but the causal facts play a role in *constituting* achievements, not causing them. This gives us reason to think that competence possession is a *sui generis* kind of property, and is not merely a species of disposition possession.

Although we fail to subsume explanation by appeal to competence under a more general kind of explanation, I actually think that this is a welcome result. Recall the question we began with, namely whether we could better understand what it is to be a performance of the kind of interest to philosophy and cognitive science by understanding it to be, or essentially depend on, competent performance. If competence possession is a *sui generis* property that only subjects (in virtue of having cognitive systems) can have, then their exercises too are *sui generis* performances that only subjects can have. Pursuing the differences between competence possession and disposition possession, as well as other stable properties that inanimate objects (like matches) can have, may then help us to understand what delineates the kinds of states and activities that essentially have subjects—that are performances—from those that do not. Exploring this line of thought in detail, however, is a task for future research.

## Chapter 3

### Competence to Know

We can see epistemology as a branch of the philosophy of mind. If we try to leave epistemology out of the philosophy of mind, we arrive at a radically impoverished conception of the nature of mind.

---

—Timothy Williamson (2000), p. 41.

#### Introduction

The Gettier Problem has shaped, and continues to shape, contemporary epistemology. It is now also beginning to shape contemporary philosophy of mind. Following Timothy Williamson (2000), several philosophers have endorsed the view that knowledge is a mental state *in its own right*—that is, not merely in virtue of its relations to belief. This central idea of knowledge-first epistemology raises new questions: What is it for knowledge to be mental in its own right? How is knowledge related to justified belief that falls short of knowledge? How might knowledge and mental states like it obtain in virtue of non-mental, non-normative facts? Knowledge-firsters unfortunately have been either silent on these issues altogether, or overtly pessimistic about the prospects of making significant progress on them.<sup>1</sup> In what follows, I provide a new argument for the claim that knowledge is a mental state in its own right and use it to motivate and develop a theory of knowledge that makes progress in answering these questions. The theory is a kind of virtue epistemology, although one importantly different from the familiar sort.

---

<sup>1</sup>Of course they accept that we can say some platitudes, such as that necessarily, if *S* knows that *p* then *S* believes that *p*. But they deny that we should be trying to explain why they hold. See, e.g., Williamson (2000) p. 31.

Virtue epistemologists claim that knowing is a kind of performance—an *apt performance* or an *achievement*—that is explained by the performer’s epistemic virtue.<sup>2</sup> Traditionally, virtue epistemologists have supposed that the virtue which explains knowledge is a competence to *believe truly*.<sup>3</sup> The view is thus an *indirect* virtue epistemology in that it aims to analyze knowledge in terms of a competence to do something other than know.<sup>4</sup> I will argue that we should reject indirect virtue epistemology in favor of a *direct* virtue epistemology, on which the competences which explain particular cases of knowledge are competences to know.

Although there is generally a pessimistic outlook on whether the Gettier Problem can be solved by a traditional theory of knowledge on which knowledge is justified true belief together with some further independent factor that is intended to rule out Gettier cases, some have claimed that indirect virtue epistemology *can* solve the Gettier Problem.<sup>5</sup> In section 1, I show that there is an in-principle problem with the indirect virtue epistemologist’s strategy, eliminating one of the last plausible explanations of Gettier cases, and so of knowledge, in terms of non-factive mental states and non-mental conditions. The argument also reveals an important mistake in indirect virtue epistemology’s account of epistemic competence and its relation to knowledge. Any account that invokes competences to do something which falls short of knowledge will fail to explain the way in which knowledge is an achievement that is due

---

<sup>2</sup>There are important questions here about how or whether one should distinguish mental states and activities, but this is not the place to discuss them. One might argue against virtue epistemology on the grounds that knowledge is a state rather than an activity. (Chrisman (2012) pursues this approach.) My own view is that *all* properly mental “states” are actually activities, because only if we conceive of the subject as engaging in mental activities can we understand why perceiving, thinking, and so on are attributable to the subject in the distinctive way that they are. However, it would take me too far afield to further discuss or defend this view. (See my Miracchi (2014c) for discussion.) In what follows I will use the term “state” without distinguishing between states and activities.

<sup>3</sup>There are two strains of virtue epistemology, reliabilist virtue epistemology, first developed by Ernest Sosa (1980), and responsibilist virtue epistemology, first developed by Linda Zagzebski (1996). Here I am only concerned with reliabilist virtue epistemology.

<sup>4</sup>Thanks to Eliot Michaelson for suggesting this terminology.

<sup>5</sup>There are other live strategies, which I cannot address here. Some maintain, for example, that justified true belief just is knowledge (e.g., Hetherington (1999, 2001); also see Weatherston (2003) for a defense of this approach), while some maintain that a “no defeaters” approach along the lines of Lehrer (1965) and Klein (1971, 1976) also solves the Gettier Problem (e.g., Lycan (2013)). I will not present arguments against these approaches in what follows, restricting my focus to indirect virtue epistemology. However, the arguments presented do apply to a number of related theories, including the “indirect” proper functionalism of Plantinga (1993) and Bergmann (2008).



to the knower's competence. If knowledge is an achievement that is explained by the knower's competence, it is explained by the knower's competence to know.

In section 2, I show how direct virtue epistemology captures the central insights of the knowledge-first approach, while making progress in answering both epistemological questions. I argue that the considerations presented in section 1 support the claim that knowledge is a mental state, and show how direct virtue epistemology makes this claim plausible and unmysterious. I also show how direct virtue epistemology can support the knowledge-first claim that knowledge is epistemically fundamental by explaining how justified belief depends both metaphysically and epistemically on knowledge, despite the necessity of justification for knowledge. I also show how direct virtue epistemology, while compatible with many views on the nature of belief, can also provide an attractive account of belief in terms of knowledge. I close by taking a step back to discuss more generally the aims and commitments of direct virtue epistemology.

### 3.1 Indirect Virtue Epistemology

According to indirect virtue epistemology, knowledge is an achievement that is explained by the knower's competence to believe truly.<sup>6</sup> There are several variations on the standard view, but this simple formulation will be adequate for our purposes:

#### **Indirect Virtue Epistemology (IVE)**

S knows that  $p$  if and only if:

- (i)  $p$  is true,
- (ii)  $S$  believes that  $p$ ,

---

<sup>6</sup>Indirect virtue epistemology is a powerful theory that is attractive for a number of reasons. See Greco & Turri (2011) for an overview. It traces back to Aristotle but in its current form was first proposed and developed by Ernest Sosa and subsequently defended in various forms by John Greco (2001; 2009; 2010; 2012), Turri (2011), and others. Other notable virtue epistemologists who hold a significant variation on virtue epistemology are John McDowell (1994), Robert Roberts and Jay Wood (2007), Christopher Kelp (2013), Duncan Pritchard (2010), and Wayne Riggs (2003). I will not discuss any of these theories here except for Pritchard's.

- (iii) S exercised a competence to believe truly in believing that  $p$ , and
- (iv) S's believing truly on the occasion in question is due to the competence to believe truly S exercised in forming the belief that  $p$ .<sup>7,8</sup>

The approach is a *belief-first* approach to knowledge, because it purports to explain knowledge in terms of true belief, a competence to believe truly, and non-mental conditions. According to indirect virtue epistemology, it is not knowledge that is characteristic of the competences that explain it, but rather true belief. The claim is then that knowledge may be specified in terms of true belief, the exercise of a competence to believe truly, and general non-mental facts relating true belief and the exercise of competence, such as facts about causal relations, modal properties, or probabilities.

In Gettier cases, a subject intuitively has a justified true belief, but fails to know. Indirect virtue epistemologists claim that a subject is justified just in case her believing is an exercise of a competence to believe truly—i.e. (ii) and (iii) of IVE are satisfied. They claim that even though subjects have justified true beliefs in Gettier cases, and so conditions (i)–(iii) are satisfied, their successes are not appropriately *due to* their competences. Thus condition (iv) is not satisfied, and the account predicts that such subjects fail to know.

Just so that we have a classic Gettier case in mind, consider what has come to be known as Gettier's Ford case.<sup>9</sup> In the case, we may suppose that Smith has a justified belief that his friend Jones owns a Ford. He picks a random proposition, *Brown is in Barcelona*, for which he has no evidence, and validly infers from *Jones owns a Ford* to *Either Jones owns a Ford or Brown is in Barcelona*. In fact, however, Jones is only renting the Ford, but Brown by chance is in Barcelona. Intuitively, Jones' belief is both true

---

<sup>7</sup>This characterization is intended to roughly match Sosa's (2007) account of animal knowledge, and Greco's (2010) account of knowledge. Some virtue epistemologists characterize the relevant virtue as a competence to *form* true beliefs (Greco (2010), p. 12), while others characterize it as a competence to believe truly. Because the latter is more general (including both competent true belief formation and retention), I work with it as the formulation here. All of the arguments presented should go through on either formulation.

<sup>8</sup>The reader may notice that not all of these conditions are independent. (iii) implies (ii) and (iv) implies (i)–(iii). I specify indirect virtue epistemology in this way only for ease of exposition.

<sup>9</sup>Gettier (1963), pp. 122–123.

and justified, but he fails to know.

Indirect virtue epistemologists claim that while Smith has a true belief that *Either Jones owns a Ford or Brown is in Barcelona*, and he exercises a competence to believe truly in believing the true proposition, he fails to know because his believing truly fails to be due to his competence. Intuitively, this is quite attractive. It seems that Smith's believing truly (as a result of *Brown is in Barcelona* being true), is not related in the right way to his exercise of competence (which is, e.g. something like a competence to believe truly by inferring *p or q* from a justified belief that *p*).

In order to provide this explanation, however, indirect virtue epistemologists must make it convincing that a subject succeeds in a way that is due to her competence to believe truly only when she knows. This, I shall argue, cannot be done. First I present a counterexample illustrating the central difficulty with the indirect virtue epistemologist's strategy; then I will present a general procedure for generating similar counterexamples.

### 3.1.1 Between a Rock and a Hard Place

Please consider the following case:

**ROCK AND HARD PLACE:** Annette is taking a walk through the countryside. Looking at what seems to her to be a sheep in the field, she forms the belief that there is a sheep in the field on the basis of her perceptual experience. In fact it is a sheepdog, but there is a sheep standing behind a rock, out of view, that the dog is keeping track of. Unbeknownst to Annette, she is in hard-working sheepdog country, in which it is very rare for a sheepdog to be in a field unless it is keeping close watch on its sheep.

This is a Gettier case.<sup>10</sup> Annette is justified, because she forms beliefs on the basis of perceptual experience (and let's assume she is typically justified when she does so), and her belief is true, but she doesn't know. Rock and Hard Place is not, however, an ordinary Gettier case. Annette is *systematically* Gettiered. This is the feature of the case that reveals the problem with the indirect virtue epistemologist's account of the epistemic competences that explain knowledge.

---

<sup>10</sup>This Gettier case is a variation on one that Roderick Chisholm gives in his (1966), p. 23.

Importantly, the case is not a kind of fake barn case.<sup>11</sup> Annette’s perceptual faculties are reliable at discriminating sheep from non-sheep in hard-working sheepdog country. It is just that, when she is in circumstances where a dog in a field looks like a sheep to her, it is highly likely that her perceptual belief that there is a sheep in the field will be true.

To show that Rock and Hard Place presents a new difficulty for indirect virtue epistemology, I will first explain how the case is a counterexample to the most plausible and well-defended proposals for solving the Gettier Problem by specifying the *due to* relation so as to exclude Gettier subjects—those of Ernest Sosa (2007, 2010) and John Greco (2012). Then I will draw some general lessons.<sup>12</sup>

### 3.1.2 Sosa (2007)

Sosa’s (2007) proposal is that a success is appropriately due to competence just in case the subject’s believing *truly*, and not just her believing, is causally explained by the subject’s competence:

Something may explain the existence of a certain entity, however, without even partially explaining why it has a given property. ... [T]he true belief that someone here owns a Ford may owe its existence to an exercise of some epistemic competence, without owing its correctness to that competence... . The Gettier victim’s belief is owed in part to his exercise of an epistemic competence without the *correctness* of that belief being similarly owed.

—Sosa (2007), pp. 95-96

Sosa’s claim here is that in Gettier cases the subject exercises her competence to believe truly, and so her exercise of competence causally explains the fact that she believes, but the fact that she believes truly (rather than falsely) is not causally explained by

---

<sup>11</sup>I will not be concerned here with fake barn cases, which are more contentious. In Goldman’s original case (Goldman (1976)) the subject, Henry, is in an area that is full of fake barn facades. If Henry had seen one of these facades, he would have (falsely) believed that it was a barn. As it happens, however, he sees a real barn and believes on the basis of his perceptual experience that it is a barn. Many have the intuition that Henry’s belief is true and justified, but that he fails to know. See Sosa (2010) for defense of the claim that subjects actually know in fake barn cases. My view is that such subjects fail even to be justified because hostile environmental conditions deprive them of the relevant epistemic competences.

<sup>12</sup>It is useful here to establish some common terminology. Thus I will use “competence” to discuss Greco’s view even though he uses “ability”. It should make no difference to the arguments presented.

her competence, but rather by some other facts. Thus her success is not (causally) due to her competence. In the Ford case, what causally explains Smith's believing *truly* has nothing to do with the exercise of competence (having instead to do with facts about what made him randomly pick *Brown is in Barcelona* rather than some other proposition). Thus, Smith's believing truly fails to be due to his competence.<sup>13</sup>

Rock and Hard Place is a counterexample to Sosa's (2007) proposal. In this case Annette's competence to believe truly that there is a sheep in the field on the basis of a perceptual experience as of there being a sheep in the field is causally responsible, not just for the existence of her belief, but also for its truth. After all, the sheepdog is only in the field because the sheep is, and so the exercise of her competence to believe truly on the basis of her perceptual experience as of there being a sheep in the field is causally dependent on the fact that makes the belief true (via perception of the sheepdog).

Furthermore, we may suppose that, had Annette not formed beliefs using her default competence to believe truly on the basis of her perceptual experiences as of there being a sheep in the field, she would have attended to subtle variations in the anatomical structure of the sheepdog (barely visible from such a distance), that would reveal that it was not a sheep. She would then have believed that there was no sheep in the field, and so would have believed falsely. Her competence thus *is* "a factor that, either singly or in combination with other factors, accounts for how the belief is true rather than false" (Sosa (2007), p. 96), and so her believing truly is causally explained by her competence.<sup>14</sup> Nevertheless, she does not know that there is a sheep in the field.

---

<sup>13</sup>There are many questions to raise about this explanation, not least of which is why it is that Smith's competence does not at least partially causally explain his believing truly. Perhaps if he had not exercised that competence, he would have exercised a different one which would have resulted in his believing falsely on the matter. Nevertheless, granting an intuitive account of causal relevance that gets the Ford case right, it is clear that the account gets other cases wrong. See Greco (2009) for other counterexamples.

<sup>14</sup>Of course, there is no uncontested inference from the obtaining of certain counterfactuals to the existence of a causal relation, but such counterfactuals seem to be quite good evidence for such a relation. The case also fits Sosa's clarification of what he means by the subject's believing truly being causally due to her competence. Absent an explanation of why the purportedly knowledge-constituting causal relation fails to obtain in this case, it should be taken to be a counterexample.

### 3.1.3 Greco (2012)

According to Greco (2012), a subject's believing truly is due to her epistemic competence (ability) just in case her competence contributes to her believing truly in a reliable way.<sup>15</sup>

A success is attributable to S's ability just in case S's ability contributes to that success in the right way, where "in the right way" means "in a way that would regularly serve relevant purposes."

—Greco (2012), p. 14.

In cases of testimonial knowledge and knowledge involving extended cognition, but not in Gettier cases, S's believing from ability contributes to S's believing the truth in the right way—i.e., in a way that would regularly serve relevant informational needs. That is, in cases of knowledge, S's abilities exploit social practices and technologies so as to produce true belief in regular, dependable ways.

—Greco (2012), p. 17.

The account explains Gettier's Ford case as follows. Because Smith was just as likely to pick a false proposition as the true *Brown is in Barcelona*, his competence fails to contribute to his believing truly in a way that would regularly produce true belief.

Rock and Hard Place, however, does satisfy Greco's conditions. It is a situation where Annette's competence reliably serves the purpose of forming true beliefs, because it is highly likely that whenever she believes that there is a sheep in the field on the basis of seeing something which looks to her to be a sheep in the field (but is in fact a dog) her belief is true (because the dog is watching a nearby sheep). Thus Rock and Hard Place satisfies Greco's conditions, but nevertheless Annette does not know.<sup>16</sup>

---

<sup>15</sup>Greco gives another account (2009; 2010) which appeals to explanatory salience in order to explain success due to competence. I do not discuss this proposal because I think there are decisive problems with it, of which Greco is now aware and which have led him to abandon the view and develop the one under consideration here. See his (2012) for discussion.

<sup>16</sup>Since Greco is trying to explain Gettier cases, and so how justified true belief may come apart from knowledge, he needs to specify a way in which the subject's competence to believe truly is involved in the acquisition of true beliefs that is distinct from the one relevant to competence possession and exercise (and thus justification). Although Greco does not explain how this type is determined, it should be clear that, on an intuitive account of the relevant way the competence is involved, the account fails to deliver correct predictions.

### 3.1.4 Sosa (2010)

Lastly, Sosa (2010, forthcoming) takes another approach to the problem that is independent of his (2007) proposal. He notes that, because on his view competences are dispositions, he may appeal to the notion of a *characteristic manifestation* of a disposition, leaving the difficult work of specifying exactly what characteristic manifestations of dispositions are to metaphysicians.<sup>17</sup> Just as a characteristic manifestation of fragility is a case of a fragile object shattering upon hitting a hard surface in the characteristic way, so a characteristic manifestation of a disposition to believe truly is a case of a believer believing truly in the way that characteristically manifests her disposition (competence) to do so. Sosa then claims that in cases of knowledge, but not in Gettier cases, the believer characteristically manifests her disposition (competence) to believe truly.<sup>18</sup>

While this approach does not explicitly specify the *due to* relation, it supposes that whatever general, non-mental facts explain what makes a success a characteristic manifestation of a disposition to believe truly will explain which cases of believing truly are cases of knowledge and which are not. The view can thus be considered a version of indirect virtue epistemology, as specified above, on which a success is due to a competence to believe truly just in case the subject's believing truly characteristically manifests her competence.

How does this approach explain why subjects fail to know in Gettier cases? Sosa (forthcoming) suggests that in Gettier cases, the subject's believing truly is not a characteristic manifestation of her disposition to believe truly because the characteristic manifestations are merely *mimicked*. As an analogy, he discusses a case where a glass is disposed to shatter upon hitting something hard, and shatters upon hitting something hard, but because a zapper who hates fragile objects hitting hard surfaces zaps and destroys the glass just upon impact. He argues that this is a case where the glass

---

<sup>17</sup>Both Sosa and Greco hold that knowledge is a success due to a competence to believe truly, where a competence is a disposition to believe truly when one believes. See e.g. Sosa (2007), p. 29, and Greco (2012), p. 18.

<sup>18</sup>See Turri (2011) for a similar proposal.

fails to manifest, and merely mimics, its disposition to shatter upon hitting something hard.

Likewise, Sosa claims that Gettier cases are also cases of mere mimicking a manifestation of a disposition to believe truly, because although the Gettiered subject believes truly, and her competence (disposition) to do so is causally involved, her believing truly fails to be related to her competence in the way characteristic of the competence. For example, in Gettier's Ford case, although Smith has a disposition (competence) to believe truly by inferring *p or q* from a justified belief that *p*, and does believe truly, his believing truly is not related to his disposition so as to characteristically manifest it (presumably because it is the truth of *q* rather than the truth of *p* that makes *p or q* true).

For Sosa's argument by analogy to be effective, the cases involving glasses that are structurally analogous to Gettier cases cannot be ones in which a glass manifests its disposition to shatter upon hitting something hard. This is so regardless of whether the glass in such a case fails to manifest *fragility*, for it is a matter of some dispute whether fragility just is being disposed to shatter upon hitting something hard (just as it is a matter of dispute whether having the competences relevant to knowledge is a matter of having dispositions to believe truly). Sosa thus cannot appeal to intuitions about whether fragility is characteristically manifested in analogous cases, but rather he must appeal to intuitions about dispositions that are more clearly similar in structure to dispositions to believe truly.

Even conceding that Sosa can in this way account for why Gettier's Ford case fails to be knowledge, the proposal cannot be extended to respond adequately to Rock and Hard Place. It is not plausible to suppose that the characteristic manifestation of Annette's disposition to believe truly is merely mimicked.<sup>19</sup> The analogous case to

---

<sup>19</sup>Of course the correct analysis of dispositions is a matter of considerable contention. Note however that the other conditions that make trouble for theories of dispositions clearly do not apply in this case. Annette's disposition to believe truly is not "reverse-cycle finked": she does not lose the disposition in conditions that are normally stimulus conditions of its manifestation, such as where a live wire is made to go dead whenever it would be touched by a conductor. Her disposition is not masked, as when a fragile object is prevented from breaking by being wrapped in packaging. These examples of troublesome cases are taken by Choi & Fara (2012) who take reverse-cycle finking, masking, and mimicking to be the three main difficulties for accounts of dispositions. They attribute the mimicry example to Lewis (1997).



consider here would be the case of an attentive zapper who does not merely by chance look at a glass and cause it to shatter upon impact, but rather takes care to ensure that if the glass were to hit something hard, it would shatter by her zapping. In such circumstances, it is quite plausible that the glass *does* have a disposition to shatter upon hitting something hard, and in a case of its hitting a hard surface and being zapped by the zapper, it would characteristically manifest this disposition. Of course, it fails to characteristically manifest a disposition to shatter upon hitting something hard *solely in virtue of the physical facts concerning the effects of impact on things of its chemical structure*, which is a disposition that genuinely fragile objects seem to have and tend to manifest upon hitting hard surfaces.

The critical question is whether, by appealing to the idea of a *characteristic manifestation* of a disposition, Sosa is in fact appealing to manifestations of much more fine-grained dispositions—dispositions that, in the epistemic case, are nothing short of dispositions to know.<sup>20</sup> In the case of our attentive zapper, there is no basis for denying that the glass manifests a disposition to shatter upon hitting something hard, although it lacks further dispositions that other things that share that disposition typically have. Likewise, there is no basis for denying that Annette manifests a disposition to believe truths on the basis of having a perceptual experience as of a sheep in the field, although she fails to manifest other dispositions normally associated with such subjects, such as a disposition to believe truths on the basis of seeing things as they are.

Thus Rock and Hard Place is a counterexample to Sosa (2010, forthcoming) as well. As far as I know, no indirect virtue epistemologist has proposed a plausible solution to the Gettier Problem which predicts that Annette fails to know. But why exactly does Rock and Hard Place pose such a difficulty for these accounts? The case is one where a bit of “bad” luck is systematically related to a bit of “good” luck, so that general properties, such as causal relevance and reliability, that indirect virtue epistemologists claim relate true belief and exercises of competence only in cases of knowledge remain

---

<sup>20</sup>The same worry arises for Turri (2011)’s proposal, which leaves “manifestation” as primitive.

in place in bad luck cases. This strongly suggests that whatever makes bad luck “bad” cannot merely be the failure of these general non-mental properties to obtain. Rather, the foregoing supports the idea that the facts constitutive of a case of knowledge in a given case are particular to the competence that explains the case of knowledge.

### 3.1.5 A General Procedure

This suggestion can now be supported by a more general argument, which parallels one presented by Linda Zagzebski (1994).<sup>21</sup> There she argues that no account on which knowledge is justified true belief together with further independent conditions can solve the Gettier Problem, by providing a procedure for inventing Gettier cases to counterexample proposals. She argues that as long as the proposal on offer allows that it is possible for the subject to have a justified false belief, one may build a Gettier case out of a case of justified false belief:

As long as there is a small degree of independence between this other element and the truth, we can construct Gettier cases by using the following procedure: start with a case of justified (or warranted) false belief. Make the element of justification (warrant) strong enough for knowledge, but make the belief false ... due to one element of bad luck. Now emend the case by adding another element of luck, only this time an element which makes the belief true after all. The second element must be independent of the element of warrant so that the degree of warrant is unchanged. ... We now have a case in which the belief is justified (warranted) in a sense strong enough for knowledge, the belief is true, but it is not knowledge.

—Zagzebski (1994), p. 69

Rock and Hard Place shows us how to specify a similar procedure for constructing Gettier cases in which the subject succeeds in the way the indirect virtue epistemologist claims is “due to” her competence to believe truly, but nevertheless fails to know. If, pace indirect virtue epistemology, what makes bad luck “bad” is that it deviates from the characteristic manifestation of the subject’s competence to know, we may start with a case of a subject exercising her competence in bad luck, and add in whatever general modal, causal, and probability relations between true belief and the exercise of a competence to believe truths that the indirect virtue epistemologist proposes suffices

---

<sup>21</sup>Thanks to Daniel Singer for suggesting that I provide such a procedure.

for the success to be due to the competence. We may thereby make it the case that the subject's believing truly, although in a bad luck situation, is "due to" her competence to believe truths. Nevertheless, she will fail to know.

*Here is our procedure:* Start with a Gettier case (perhaps obtained by Zagzebski's procedure). Then specify a further fact that is independent of the justification (warrant) that makes the elements of bad and good luck systematically related, so that if the subject were to form beliefs in the way that yields justification *and encounters the bad luck*, she would encounter the good luck as well. If further general facts involving causation or probabilities are involved, one may specify general causal or probabilistic relations between the bad and good luck that ensure these facts too. This will suffice to make it that, in a case where the subject exercises her competence and both the bad and good luck obtain, she will succeed in the way that the indirect virtue epistemologist claims makes her success due to her competence to believe truly. Nevertheless she will fail to know.

The procedure leaves undisturbed the causal regularities that purportedly constitute knowledge in normal cases, and establishes regularities with exactly the same properties in bad luck cases. In doing so, it exploits two central properties of the indirect virtue epistemologist's account: that competences to believe truly may be fallible and that the "due to" relation is specified in competence-general terms. To see an example of how the procedure works, consider the following case:

**DOUBLE PLAY:** Creola has two friends, Fred and George. For some reason, Fred lies to George and tells him that he did not see a play that evening, when he actually has. By chance George, trying to deceive Creola, tells her that Fred *did* see a play that evening. Creola, suspecting nothing, forms a true belief that Fred saw a play that evening.

This is a Gettier case. Creola is justified, because she forms beliefs on the basis of testimony (and let's assume she is typically justified when she does so and has no reason to think George is unreliable or uncooperative on such matters), and her beliefs are true, but she doesn't know. We may now use our procedure on this case to devise a systematic Gettier case that counterexamples all three proposals above. What we

must do is emend the case so that the bad luck (encountering a deceiver) is systematically counteracted by good luck (the deceiver having misleading information) in a way that makes Creola reliably form true beliefs. We leave general facts relating her competence to believe truly on the basis of testimony intact—e.g., that she normally believes truly because she encounters reliable and cooperative informants telling the truth. Here is our resulting case:

**DOUBLE TROUBLE:** Creola's two friends Fred and George sometimes like to play tricks. Fred goes on a month-long trip, and communicates with George daily. Fred decides that he will play a trick on George and tell him the complete opposite of what actually happens to him on his trip. George, charged with relaying news of Fred to Creola, decides he will play a trick on her, and tells her the complete opposite of what Fred tells him. As it so happens, Fred and George's subterfuges reliably cancel each other out, so that Creola reliably receives true information about Fred's trip. She forms beliefs on the basis of George's testimony just as she would in any other standard case of belief-formation on the basis of testimony.

Creola is systematically Gettiered, just as Annette is in *Rock and Hard Place*. We may suppose that, had Creola not formed beliefs using her default competence to believe truly on the basis of testimony, she would have attended to subtle signs in George's voice, eyes, etc., that would reveal his subversion. She thus would have believed he was lying, and have formed false beliefs about Fred's trip. Thus her believing truly, and not just her believing, is causally explained by her competence, *pace* Sosa (2007).

Moreover, because Fred and George's uncooperativeness systematically cancel each other out, Creola's exercises of competence to believe truths on the basis of testimony *do* contribute to her successes in a way that regularly serves her epistemic purposes (of forming true beliefs). When she forms beliefs on the basis of George's testimony, she forms beliefs in a way that is highly likely to produce true beliefs. Thus *Double Trouble* is a counterexample to Greco's (2012) proposal.

Lastly, there is no basis for denying that Creola manifests a disposition to believe truths on the basis of testimony, although she fails to manifest other dispositions such subjects normally do, such as a disposition to believe truths on the basis of a reliable and cooperative informant telling the truth. Thus *Double Trouble* is a counterexample to Sosa (2010, forthcoming) as well.

### 3.1.6 Safety to the Rescue?

The reader might suppose that the difficulties raised above for indirect virtue epistemology can be dealt with by appeal to a distinct anti-luck condition, for example a safety condition, as Duncan Pritchard (2010) proposes. However, our procedure works equally well with Pritchard's proposal as it does with those above. According to the procedure, we start with a Gettier case, and then emend the case so that the proposed knowledge-ensuring condition holds by establishing a connection between the "bad" luck and "good" luck conditions.

Let us start with the Double Trouble case. We then establish a connection between Fred's trick playing (the bad luck) and George's trick playing (the good luck) so that in all of the closest possible worlds where Fred plays his trick George does too. Suppose that Fred and George have gotten their respective ideas to play their tricks when they watched a movie together (say, in which one friend plays a prank on another), and that neither of them would easily have seen the movie by himself or have gotten the idea for his prank in another way. If this is so, then it is quite plausible that in the closest possible worlds where Creola forms beliefs on the basis of testimony, she forms true beliefs. She will either be in a world where both Fred and George are playing tricks, and so her belief is true, or she will be in a world where neither is playing tricks, and so again her belief is true. Her beliefs in the emended case are thus safe, true, and a matter of her exercising a competence to believe truly. Nevertheless she does not know.<sup>22</sup>

The procedure presented shows that there is no general relation between an exercise of a competence to believe truly and believing truly that is constitutive of knowledge. If any such condition is proposed, we will be able to use the procedure to devise a case in which it obtains but the subject intuitively fails to know. The ability to do this no matter what sort of competence-general condition is imposed shows that the

---

<sup>22</sup>It should be clear that the difficulty for indirect virtue epistemology posed by Double Trouble does not depend in some way on other purported difficulties for virtue epistemology regarding testimony and credit (see e.g. Lackey (2007, 2009)). Double Trouble is the straightforward result of applying our procedure to Double Play. The rules of the procedure for developing systematic Gettier cases have nothing to do with the particularities of knowledge on the basis of testimony. (Thanks to Jessica Brown for suggesting that I discuss this.)

virtue epistemologist cannot rule out knowledge-precluding bad luck by appeal to competence-general facts. Whatever the particular truth-ensuring facts are that are constitutive of knowledge, they also determine the natures of the competences that explain particular cases of knowledge. If knowledge is an achievement that is explained by the knower's epistemic competence, the epistemic competence cannot be anything less than a competence to know.

### **3.2 Direct Virtue Epistemology**

I shall from here on assume that knowledge is an achievement that is to be explained by the knower's competence to know, and will argue that cases of knowledge are *exercises* of competences to know, thus establishing the central claim of direct virtue epistemology. I then show how this claim explains how knowledge is a mental state in its own right, and impose some necessary conditions on competences to know that explain the relationship between knowledge and justification. Finally, I consider how direct virtue epistemology may account for the relationship between knowledge and belief.

#### **3.2.1 Knowledge Is an Exercise of a Competence to Know**

The claim that knowledge is an exercise of a competence to know is distinct from the claim that knowledge is an achievement that is due to a competence to know. This is made clear by comparison with indirect virtue epistemology. According to indirect virtue epistemology, the competences which explain knowledge are competences to believe truly, but the relevant exercises of such competences are mere believings, which may or may not be true. This is not an idle feature of the theory: that exercises of competence are neutral with respect to success or failure is crucial to the indirect virtue epistemologist's explanation of how a subject may have a justified false belief. This occurs just in case the agent exercises her competence to believe truly, but through some unlucky event happens to believe falsely.

The same considerations that force rejection of the claim that knowledge is a success due to a competence to believe truly also require us to reject the claim that knowing is a matter of exercising a competence to know together with the obtaining of further non-mental conditions. The problem for indirect virtue epistemology above was that what made a case have knowledge-precluding bad luck was particular to the competence at issue, and could not be specified generally. We cannot understand what it is for knowledge to be an achievement that is due to the knower's competence by specifying a non-factive exercise (believing), an independently specifiable success condition (believing truly) and a general non-mental relation between the non-factive exercise and the success.

This difficulty is not resolved if we suppose the success condition to be knowing, and maintain that the exercises of competences to know can obtain either in cases of knowledge or in cases that fall short of it. If this were true, then what it is to know would be specifiable independently of the particular competence in question, and there would have to be a general relation that holds between cases of knowledge and exercises of competences to know. However, it should be clear that the procedure for generating systematic Gettier cases would be equally applicable to a proposal along these lines. Whatever the general relation between the exercise of competence and the state that is purportedly knowledge, we will be able to find a case where that relation obtains but the subject's being in such a state fails to be due to her competence in the way that seems distinctive of knowledge. What it is for something to *be* a case of knowledge is not independent of the particular competence that brings it about.

If the considerations I have presented so far are compelling, then indirect virtue epistemologists have not only been wrong about which epistemic competences explain knowledge, but they have also had the wrong account of *competence*. Competences to know have as exercises the very achievements they are invoked to explain—i.e. particular cases of knowledge. These particular cases of knowledge are such *because* they are exercises of such competences.

### 3.2.2 Knowledge Is a Mental State

Knowledge, then, is an exercise of competence. This is all we need to establish and clarify the claim that knowledge is a mental state in its own right, at least in the core sense at issue between knowledge-firsters and belief-firsters.

I have so far left the claim that knowledge is a mental state at a fairly intuitive level. Now it is important to become more precise. Williamson claims that “knowing is a mental state in *every* reasonable sense of the term” (Williamson (2000), p. 28). There seems to me to be at least one reasonable sense of “mental” in which it is used to distinguish certain person-level states or activities from ones that involve the agent’s body and surroundings.<sup>23</sup> For example, one might ask a child whether she can do long division mentally, or “in her head”, meaning without the use of pencil and paper.

This distinction, while clearly reasonable, seems to me equally clearly *not* to be what the knowledge-firster is after when she claims that knowledge is a mental state. After all, dividing 5899 by 17 with pencil and paper is a perfectly good way of coming to know that  $17 \times 347 = 5899$ .<sup>24</sup> What does seem to be the crucial claim the knowledge-firster is after is that knowledge is *personal*, whether or not purely mental in this sense. What is at issue is whether the knower is the subject of her knowing merely in virtue of being the subject of her believing or justifiedly believing. This, too, is a perfectly reasonable sense of “mental”, and the knowledge-firster is right. The knower is the subject of her knowing because of the kind of state it is, not merely in virtue of its being constituted by other mental states. This is because knowing just is an exercise of competence, and exercises of competences possessed by persons are personal states and activities.

It is not contentious that *exercises* of competences (capacities/ abilities/ methods)

---

<sup>23</sup>Thanks to an anonymous reviewer for pushing me to clarify this point.

<sup>24</sup>Williamson (2000) notes this usage (see Ch. 2). I think we merely differ on whether or not this sense of “mental” is *reasonable*. Note that this construal of “mental” will also likely classify perceptual tracking (which constitutively involves ocular movements) and many other intuitively “mental” states and activities as non-mental. This seems to me to be harmless terminology as long as the stipulated meanings are kept in mind.



are mental (in the sense clarified). The idea that exercises of competences (capacities/ abilities/ methods) are for that reason personal states and activities is implicit in much discussion in philosophy of perception, language, and action. Throughout each of these domains, it is assumed that perceptual capacities yield perceptual experience, that linguistic competence yields language comprehension and production, that knowledge-how yields basic intentional actions.<sup>25</sup> Tyler Burge, in discussing perceptual capacities, even goes so far as to claim that: “Perceptual states are realizations of individuals’ capacities. I think that this claim is a priori” (Burge (2010), p. 369). Indeed, this idea is part of the strategy of the indirect virtue epistemologist, who claims that what it is to believe justifiedly is to exercise a competence to believe truly. According to this view, the mentality of *justification* is not in dispute; only the mentality of knowledge.

What *is* contentious is whether the achievements that are explained by such competences are also mental in their own right—in the case at hand, whether the knower is the subject of her knowing because of the kind of state it is, or merely because it is constituted by another state she is the subject of (believing or justifiedly believing). If indirect virtue epistemology were adequate, then doubts about the mentality of knowledge would be vindicated. We would be able to specify the kind of state knowledge is in a way that only appeals to or makes plausible the mentality of belief, justified belief, and competences to believe truly.

But as I have shown above, indirect virtue epistemology cannot adequately specify the kind of state knowledge is. In order to characterize the way in which knowledge is an achievement that is due to the knower’s epistemic competence, we must suppose that knowledge itself *is* an exercise of competence. This is a way of characterizing the nature of knowledge on which it is a mental state in its own right, not merely in virtue of being constituted by other mental states (if it in fact is).

Rather than being an inductive argument based on the failure of the post-Gettier

---

<sup>25</sup>As a somewhat random sampling of recent work illustrating this point, see Burge (2010), Devitt (2006), Dickie (2012), and Schellenberg (2013).

literature, or an argument from the explanatory or conceptual centrality of knowledge, the argument presented here is a new, much more straightforward argument for the claim that knowledge is a mental state, based on reflections about the nature of knowledge and the features which plausibly explain it. Beginning with the antecedently plausible claim that knowledge is an achievement that is to be explained by appeal to competences of the knower, and by investigating the relationship between knowledge and the kind of epistemic competence that explains it, we have reached the result that knowledge just is an exercise of a competence to know, and thus is mental in its own right.<sup>26</sup>

The argument above not only establishes that knowledge is a mental state, it also makes the claim non-mysterious. The view that knowledge is a mental state because it is an exercise of competence puts it in a good deal of company, helping us to understand it by subsuming it under this general kind and comparing it with other states and activities that are exercises of competence. It also does not motivate pessimism about ultimately providing a complete metaphysical analysis of knowledge. Indeed, we have no more reason to be pessimistic in this regard than we have to be pessimistic about providing analyses of any other mental states or activities that are exercises of competence.<sup>27</sup>

### 3.2.3 Knowledge and Justification

Can direct virtue epistemology explain justification in terms of knowledge, and so support the knowledge-first thesis that knowledge is epistemically fundamental? It seems quite plausible that there are beliefs that are justified to the same degree as cases of knowledge, but which fail to be knowledge. This may occur because through some

---

<sup>26</sup>While my claim that knowledge is a mental state is more narrow than Williamson's, the considerations presented in favor of it lead to a conclusion that is far more general. For those who think Gettier cases are species of the larger class of *deviant causal chain cases*, the foregoing suggests the plausibility of an *achievement-first* virtue-theoretic study of achievements, in any domain. This, I think, is exactly right, and is the subject of my Miracchi (2014a).

<sup>27</sup>I think we actually have good reason to be optimistic. Frances Egan (2013), among others, has proposed that computational explanations are explanations of the processes that underlie competent performance.) If that is so, then direct virtue epistemology brings knowledge cleanly within the domain of phenomena explicable by a flourishing scientific research program.

unlucky event a subject believes falsely or is Gettiered. Explaining how there could be such *merely* justified beliefs, and how there could be knowledge-level justification (henceforth “justification”) that cases of knowledge and cases of non-knowledge can share, has been a sore point for knowledge-first epistemology.<sup>28</sup> There are two difficulties here.

First, justification seems to be necessary for knowledge, but not the reverse. How is this so, if the knowledge-firster is right that knowledge is metaphysically prior to all other epistemic states? Usually we take necessity to be a guide to metaphysical priority. If we should not do so here, how else might these modal facts be explained? Second, justified beliefs that fall short of knowledge—whether true or false—seem to be cases in which the subject believes well, and in much the same way she does when she knows. However, if neither the metaphysical nor the epistemic status of knowledge is to be explained in terms of justification, how might knowledge have this positive epistemic status in common with states that fall short of it?

Indirect virtue epistemology sets a high bar for explaining these two facts. According to it, a belief is justified just in case the subject exercises a competence to believe truly in so believing. Further facts are required for knowledge, but what explains the competence, or skill, of the beliefs is the same. Thus indirect virtue epistemology can explain both the necessity of justification for knowledge (and not the reverse) as well as their common positive epistemic status *qua* exercise of epistemic competence.<sup>29</sup>

Direct virtue epistemology can incorporate the insight of indirect virtue epistemology that it is facts about common epistemic competences that explain both the modal relationship between knowledge and justification, as well as the common positive epistemic status they share. In order to do so, however, it must take competences to be different from both of the two main existing approaches. On the first approach,

---

<sup>28</sup>Although Williamson (2000) says a good deal about the nature of evidence (in defense of the claim that one’s evidence just is what one knows), he says surprisingly little about knowledge-level justification. Moreover, by imposing a certain kind of method-relative sensitivity requirement on knowledge on which, necessarily, if *S* knows that *p* via method *M*, then in *all* nearby possible worlds if *p* is false then does not believe *p* via *M* (p. 154-5), he precludes the kind of account developed below.

<sup>29</sup>To use Sosa’s terminology, a shot may be *adroit*, even if it fails to be *apt*, and the same competence can be responsible for both kinds of shots. See, e.g., Sosa (2007), p. 22.

competences (capacities/abilities/methods) are fallible and their exercises are non-factive. On the second, competences are infallible and their exercises are factive.<sup>30</sup> Instead of supposing either that justified belief is the neutral exercise of competence, which may or may not amount to knowledge depending on circumstances, or that exercises of epistemic competence are always achievements (knowledge), we need an account on which competences may have two kinds of exercise: one of which is constitutively an achievement and the other of which is constitutively a failure.<sup>31</sup> That is, we need to understand how some exercises of competences to know—the ones that determine what the competence is a competence to do—just are cases of knowledge, while Gettier cases and cases of justified false belief are degenerate exercises of the same competence. I will impose some (merely!) necessary conditions, to show how this is possible.

To begin, let us suppose that we are placing conditions on a competence  $C_K$  to know. Competences to know may be competences to know on the basis of testimony, on the basis of deductive or inductive inference, on the basis of perception, on the basis of memory, or any other way in which one may come to know, or continue to know, that  $p$ . We may associate with such a competence a way of knowing  $W_{CK}$ , which is its characteristic manifestation. For example, the competence to know on the basis of testimony has knowing on the basis of testimony as its characteristic manifestation.<sup>32</sup> We may then specify the *manifestation conditions* of  $C_K$ , which are the conditions that in a particular case (against a background of possession of  $C_K$ ) constitute knowledge in way  $W_{CK}$ :

---

<sup>30</sup>For proponents of the first approach, see, e.g., Bergmann (2008); Burge (2010, 2011); Greco (2010, 2012); Plantinga (1993); Schellenberg (2013); Sosa (2007). For proponents of the second, see, e.g., Millar (2008); Williamson (2000). Williamson takes this approach with respect to methods. The safety condition he puts on knowledge requires methods that are only used to acquire knowledge.

<sup>31</sup>It is often supposed that if a competence or capacity is fallible, then there must be a non-factive mental state in common to the two cases. See Burge (2005) for defense of this claim. Here I show that this is not so for the case of knowledge. For more general arguments to this effect, see Miracchi (2014a).

<sup>32</sup>Probably competences to know will be much more fine-grained, but I simplify here for the sake of exposition.

**Manifestation Conditions:** The manifestation conditions of  $C_K$  are whatever operations of subpersonal cognitive mechanisms and external conditions together (against a background of possession of  $C_K$ ) constitute a particular case of knowing that  $p$  in the way characteristic of the competence ( $W_{CK}$ ).

A competence to know is **manifested** just in case its manifestation conditions obtain.

Clearly, this condition is not very informative, and although it is not stated here the manifestation conditions will of course involve facts that ensure the truth of the doxastic state constituted by the manifestation conditions. The point of specifying the manifestation conditions is not to illuminate *them*, but to explain how there may be fallible competences *to know*, and so to explain the relationship between knowledge and justification. (Note also that since on direct virtue epistemology knowledge is, as a matter of its nature, an exercise of competence, the obtaining of the manifestation conditions will only constitute a particular case of knowledge against a background of possession of the relevant competence. Thus the qualifications above.)

Of the manifestation conditions, we may isolate just the subpersonal cognitive mechanisms whose operations in a particular case partially constitute the subject's knowing that  $p$  in way  $W_{CK}$ :

**Basis Condition:** The basis of  $C_K$  is fully constituted by the subpersonal cognitive mechanisms of the subject  $S$  whose operations partially constitute  $S$ 's knowing that  $p$  in way  $W_{CK}$ .

I may also remain non-committal about the nature of these mechanisms and their operations.<sup>33</sup> Just distinguishing the basis of a competence to know from its manifestation conditions as above allows us to impose a different kind of reliability condition

---

<sup>33</sup>The operations of these mechanisms typically will have functional, and perhaps computational, specifications. I am using the word "cognitive" in a broad way to include, e.g., perceptual mechanisms and other mechanisms that are the object of study for cognitive science. I here leave it open whether some bodily, and not just brain mechanisms are cognitive mechanisms.

than is normally imposed, and so to explain how competences to know may be fallible. Instead of supposing that the central condition on epistemic competences is that their exercises must be likely to produce true beliefs, we may suppose that it must be sufficiently likely that whenever the subpersonal cognitive mechanisms that constitute the basis of the competence are operative, the conditions constitutive of knowledge also obtain:

**Proficiency Condition:** The proficiency condition of  $C_K$  requires that the objective probability of the manifestation conditions obtaining conditional on the basis of the competence being operative be sufficiently high. I.e.,  $Pr(M|O_B) \geq n$ , for some sufficiently high  $n \in (0, 1]$ .<sup>34,35</sup>

According to the present proposal, the exercises (constituted by the manifestation conditions) that are fundamental to competences to know are cases of knowledge, not anything that falls short of it. Thus there is no room for a success condition (e.g. true belief) distinct from the achievement condition (knowledge), and so no room for Gettier cases. The proposal also explains how the claims I argued for above about the nature of knowledge and its relation to competence may be correct. I argued that in order to explain how knowledge is an achievement that is due to competence, it must not only be the case that a certain way of knowing is characteristic of the competence (section 1), but it must also be true that the nature of a way of knowing is not independent of the competence which explains particular cases of knowing in that way (section 2.1).

---

<sup>34</sup>For ease of exposition, I am representing the threshold as a real number  $n \in (0, 1]$ , but I need not make any commitments about how close to 1  $n$  must be, whether the threshold for the epistemic domain is the same as for other domains (e.g. baseball), or whether the threshold is determinate. (I think it is probably indeterminate.) I also need not make any commitments about whether the threshold is context-sensitive. E.g., if subject-sensitive invariantism (Stanley (2005)) is true, then  $n$  will vary with practical interests.

<sup>35</sup>While I need not commit to a particular view of probability here, I am supposing that appeal to objective conditional probabilities is less contentious than appeal to objective unconditional probabilities (see Hájek (2007) for defense of this claim), and I am supposing that some kind of non-frequentist realist account of them is in the offing. I am also supposing that such probabilities are true at particular times. Changes in, e.g., causal regularities over time might result in either the acquisition or loss of competences to know by changing whether or not the proficiency condition is met.

The above conditions show how both of these claims may be true. It is only because the manifestation conditions are related to the basis of the competence in the way specified by the proficiency condition that the manifestation conditions constitute particular cases of knowledge.<sup>36</sup> (Recall the qualifications about manifestation conditions constituting knowledge only against the background of competence possession.) Thus a way of knowing only is such because it is an exercise of competence. Nevertheless, for any competence to know, a certain way of knowing is characteristic of the competence because the proficiency condition requires that the basis of the competence be operative, for the most part, only in the manifestation conditions—i.e. only in the conditions that on a particular occasion are sufficient for knowing in a certain way.<sup>37</sup>

Thus the proposed necessary conditions provide a plausible way of clarifying and precisifying the claims made about knowledge and competences to know thus far. Now I shall show that they also enable us to explain how competences to know may be fallible. Because the proficiency condition does not require that the probability of the manifestation conditions obtaining conditional on the basis being operative be 1, it is in general possible for the basis to be operative in cases where not all of the manifestation conditions obtain. This allows for the possibility of degenerate exercises, *deviations* from manifestation conditions.

Let us suppose that when the basis of a competence to know is operative in conditions that *deviate* from the manifestation conditions it is *degenerately exercised*:

A competence to know is **degenerately exercised** just in case its basis is operative, but not all of the manifestation conditions obtain.

---

<sup>36</sup>This is true even if it turns out that the proficiency condition is not sufficient for establishing competences to know. I remain neutral here on whether, once the manifestation conditions and basis of a competence are fully specified, the conditions articulated above are sufficient for possession and manifestation of competences to know, or whether probability is too weak a notion to play the needed role here. However, I should note that it is crucial to the project that competences not turn out to be so sophisticated as to themselves require propositional knowledge, as Stanley (2011) claims skills do. Then the account would be viciously circular. If competences are dispositions as Sosa and Greco claim, then this worry is avoided.

<sup>37</sup>Note that this condition need hold only “for the most part” in order to allow that competences may be fallible.

Finally, we may suppose that if either of these cases obtain, the competence is *exercised*:

A competence to know is **exercised** just in case it is either manifested or degenerately exercised.

We can now more precisely state direct virtue epistemology:

### Direct Virtue Epistemology

- (i) All competences to know have bases, manifestation conditions, and a proficiency condition relating them as specified above.
- (ii) For any case of **knowledge**  $k$ ,  $k$  is a manifestation of some competence  $C_K$  to know.
- (iii) For any case of **mere justified belief**  $j_m$ ,  $j_m$  is a degenerate exercise of some competence  $C_K$  to know.
- (iv) For any case of **justified belief**  $j_n$ ,  $j_n$  is an exercise of some competence  $C_K$  to know (i.e. it is either a manifestation or a degenerate exercise of some competence  $C_K$  to know).

How then does the direct virtue epistemologist explain justified belief in terms of competences to know, so that justification is necessary for knowledge, is an epistemic status in common to cases of knowledge and cases that fall short of knowledge, and is nevertheless metaphysically and epistemically posterior to knowledge?

Justified belief, on this approach, is a disjunctive kind, whose disjuncts are knowledge and mere justified belief. Since a disjunctive kind is, necessarily, instantiated just in case one of its disjuncts is, it is impossible for there to be a case of knowledge without a case of justified belief and thus justified belief is necessary for knowledge.

However, since a disjunctive kind is metaphysically posterior to both of its disjuncts, justification is metaphysically posterior to both knowledge and to *mere* justified belief.<sup>38</sup> Mere justified beliefs, moreover, are themselves metaphysically posterior

---

<sup>38</sup>Those familiar with the debate on perception between disjunctivists and non-disjunctivists will see applications of this proposal to perceptual experience. I am currently developing such an account.



to knowledge. They are degenerate exercises of competences to know, and so are not neutral mental states or activities which are specifiable independently of the obtaining of external conditions. They are *constitutively failures*. They are dependent for their status as exercises of competence on the fact that they are partially constituted by the operation of the bases of competences to know, and that such operations obtain in conditions that deviate from the conditions constitutive of knowledge. Justification, then, is a disjunctive kind whose disjuncts are knowledge and a state metaphysically dependent on knowledge, and so is entirely metaphysically dependent on knowledge.

Lastly, direct virtue epistemology can explain how knowledge and mere justified belief have a positive epistemic status in common, one dependent on that of knowledge. As exercises of the same competence to know, a case of knowledge and a case of justified belief are cases of belief that are in an important sense equally likely to be cases of knowledge. Although the proficiency condition governs the facts constitutive of manifestations and degenerate exercises of competences to know, it establishes a further conditional probability: the probability that a belief will be knowledge given that it was an exercise of the competence to know will also be equal to  $n$ , and so will be quite high. Exercises of a competence to know, then, whether cases of knowledge or cases of justified belief that fall short of knowledge, are *qua exercises of competences to know* likely to be cases of knowledge. Moreover, cases of knowledge and cases of merely justified belief that are exercises of the same competence to know will be *equally likely* to be cases of knowledge (*qua exercises of that competence*). Thus, insofar as reliability with respect to an epistemic good is itself an epistemic good, all justified beliefs share a common positive epistemic status, one that derives from its relation to knowledge—reliability with respect to knowledge.<sup>39</sup>

---

<sup>39</sup>There are further positive features of direct virtue epistemology's account of justification which I do not have room to discuss here. For example, like indirect virtue epistemology, it can explain the positive epistemic status of justified belief without incurring a regress problem. It navigates between the "raft" and the "pyramid" by appeal to competences, for which the question of justification does not arise (see, e.g., Sosa (1980)). Another upshot of direct virtue epistemology is that, again like indirect virtue epistemology, it can explain how it is possible for a person who fails to know to be *better justified* than someone who knows. A subject's degenerate exercise of competence may nevertheless be an exercise of a much more reliable competence than that of another person who manifests her competence. In this case we may still say that the former is better justified than the latter because the competence she exercises is more likely to produce knowledge.

Direct virtue epistemology thus invokes competences to accomplish two central tasks for a knowledge-first virtue epistemology. It explains what it is for knowledge to be a mental state in its own right, by explaining how knowledge is an exercise of competence. Moreover, it explains what justification is and why it is necessary for knowledge but nevertheless metaphysically and epistemically depends on knowledge. Moreover, it accomplishes both of these tasks appealing to quite commonplace notions—probabilistic facts and subpersonal cognitive mechanisms. It just puts them to use in a new way.

### 3.2.4 Mere Belief

I have said little about belief as such. First, that knowledge is a mental state in its own right does not show that knowledge is not partially constituted by belief. Consider a related example. I know how to get just about anywhere I have the address of in the U.S.A., in part because I know how to look up directions on Google. My knowledge of how to do the former depends on my knowledge of how to do the latter. In cases where I achieve the former, my doing so is partially constituted by my achieving the latter. (It is not, e.g., merely an enabling condition.) The claim that getting where I wish to go is something that I do is not threatened just because it is partially constituted by other things I do. Although there are useful distinctions to draw between *knowledge how* and competence more generally, the same applies to our case. What if knowledge, albeit a mental state in its own right, is partially constituted by belief, and having a competence to know is partially constituted by having a competence to believe? Gettier cases do not weigh in on this question.

However, insofar as epistemic norms seem to be applicable to beliefs as such (in ways they are not to, e.g., desires, hopes, etc.), it is plausible that belief should be explained in terms of knowledge. While it is not clear to me exactly how this might go, let me close with a proposal. Some have claimed that belief constitutively aims at knowledge—that belief just is the sort of commitment to  $p$  that aims at knowing.<sup>40,41</sup>

---

<sup>40</sup>See, e.g., Bird (2007); Sutton (2007).

<sup>41</sup>It is worth pointing out that direct virtue epistemology as defined in section 2 is neutral on the

Might this proposal allow the direct virtue epistemologist to explain belief in terms of knowledge?

The account of competences to know presented here may adopt a further feature that is not entailed by the preceding but is a natural development of it, and is useful for our purposes here. While there is much disagreement about the place of the intentional in the natural world, one strong line of thought supposes that causal regularities between a person's cognitive system and her environment play a central role.<sup>42</sup> It is thus not implausible to suppose that the very same facts that make it the case that manifestation conditions are achievements and degenerate exercises are failures to achieve also make it the case that in exercising her competence the agent *aims* to achieve what she either achieves or fails to achieve.

While not providing a full account here, I am suggesting that the causal regularities constitutive of competences and their exercises are also constitutive of the agent's performing with the aim to achieve when she exercises her competence. This would of course be a kind of performing with an aim that would not require the acting agent to be able to *reflect* on her activity, or to *endorse* her activity in any way beyond the performing of it. It is even compatible with reflectively *rejecting* the activity, as an addicted smoker still counts as lighting up a cigarette with the aim of doing so even though her reflectively endorsed aim is to quit.

It would, however, suffice to make cases of knowledge and mere justified belief performances that are properly activities of an agent, rather than something she or parts of her are merely causally involved in. Since we are in the realm of some of the most basic personal performances, we should be looking for the facts in virtue of which they are performances, in virtue of which agents *do* them, rather than be merely causally involved in their occurrence. It seems that looking at the causal facts

---

question of whether knowledge even entails belief. Given that there is some contention about whether this is so, this is all the better for direct virtue epistemology. For recent discussion in experimental philosophy, see e.g., Myers-Schulz & Schwitzgebel (2013), Buckwalter et al. (2013), and Rose and Schaffer (forthcoming). For earlier challenges to the claim that knowledge entails belief, see Radford (1966) and Black (1971). My view is that philosophers' discussions of belief do not pick out a unique kind, but rather many interesting phenomena worthy of study.

<sup>42</sup>E.g., Fodor (1990a); Dretske (1981, 1986).

constitutive of competence possession and exercise is not a bad place to start.

Moreover, by adopting this extension of the account, the direct virtue epistemologist can explain why knowledge and justified belief entail belief, a desideratum which is left completely unexplained by Williamson (2000), but which falls out as a natural consequence of the current proposal. If belief just is the kind of commitment that constitutively aims at knowledge, and knowledge and mere justified belief constitutively aim at knowing *qua* exercises of competences to know, then they constitutively aim at knowing, and thus these exercises *are* beliefs.

If belief just is the kind of commitment that constitutively aims at knowledge, and knowledge and mere justified belief constitutively aim at knowing *qua* exercises of competences to know, then the direct virtue epistemologist is two-thirds of the way there to explaining belief in terms of knowledge. But what about unjustified belief?

While being an exercise of a competence to know is one way of aiming at knowledge, there is no reason to think that it is the only way. Indeed, aiming at knowing by doing other things one hopes will result in knowing is quite common. We may suppose then, that rather than there being a deep explanation for why *belief* aims at knowledge (of the sort we have for why knowledge and justified belief do), that belief just is commitment to *p* that aims at knowing, however that aim may arise.

Mere unjustified belief, then, instead of being the neutral foundation of justified belief and knowledge, may be the most degenerate case of the three. *It* earns its status as a doxastic state only in virtue of the relations it bears to knowledge, and this is a connection at a remove—it merely aims at knowledge.

### 3.3 Conclusion: What Kind of Theory is Direct Virtue Epistemology?

I have argued that no indirect virtue epistemology can solve the Gettier Problem, and thus removed one of the remaining plausible accounts of knowledge in terms of non-factive mental states and non-mental conditions. I then argued that we should instead adopt *direct virtue epistemology*, on which knowledge is a manifestation—a characteristic exercise—of a competence to know. But what kind of theory is direct virtue

epistemology? Isn't the account circular? Why is it not viciously so? How exactly is the account intended?<sup>43</sup>

Direct virtue epistemology accomplishes three tasks: First, it provides an informative description of the nature of knowledge and an account of epistemic normativity. Second, it provides an informative *partial* metaphysical analysis of knowledge. Third, it provides a *full* metaphysical analysis of justification in terms of knowledge and competences to know. In closing, I discuss each of these in turn.

### 3.3.1 The Nature and Normativity of Knowledge

Although here I have focused on the differences between direct virtue epistemology and indirect virtue epistemology, they are both kinds of virtue epistemology, and as such share central features in common. Most importantly, both theories claim that knowledge is an achievement that is due to the knower's competence. This is a positive description of the nature of knowledge which predicts that we can better understand epistemic properties by understanding features of competent performance more generally.<sup>44</sup>

This view about the nature of knowledge allows for a natural virtue-theoretic account of the positive epistemic status of knowledge. Epistemic normativity is a kind of performance normativity; knowledge enjoys its distinctive epistemic status precisely because it is an achievement that is due to competence.<sup>45</sup> Moreover, because direct virtue epistemology claims that knowledge is an exercise of competence, it also has the positive epistemic status of *being* a competent performance (just as justifiedly believing does on indirect virtue epistemology).<sup>46</sup>

---

<sup>43</sup>Thanks to an anonymous reviewer for suggesting I address this question explicitly.

<sup>44</sup>Direct virtue epistemology of course disagrees with indirect virtue epistemology about the structure of epistemic competences and how they explain knowledge. This disagreement thus suggests a much more general disagreement about the structure of competences in general and how they explain the achievements that they are invoked to explain. Whereas the traditional approach aims to explain achievements in terms of indirect competences—competences to succeed where success falls short of the achievement in question—the foregoing suggests that achievements in general must be explained by competences to achieve, which have two kinds of exercise, those that are constitutively achievements and those that are constitutively failures.

<sup>45</sup>See Greco (2010), pp. 42-46 for further discussion.

<sup>46</sup>Note that justified false beliefs and Gettier cases also have the same positive epistemic status for the

This way of explaining the nature and normative status of knowledge is not solely available to indirect virtue epistemology, but is rather a feature of virtue epistemology more generally. This is because the form of explanation is not that of reducing the explanandum to prior and independent phenomena, but rather of subsuming the explanandum under more general kinds that we have some independent understanding of—achievements due to competence, exercises of competence, and performance normativity.

Direct virtue epistemology takes this approach one step further than indirect virtue epistemology and appeals to facts about exercises of competence to explain what it is for knowledge to be a mental state (section 2.2). Knowledge is a mental state because it is an exercise of a competence of the subject, and the exercises of such competences are mental. Knowledge is thus a mental state because of the kind of state it is, not merely in virtue of being constituted by other mental states or activities.

Thus direct virtue epistemology provides an informative account of the nature of knowledge and epistemic normativity while maintaining that they cannot be analyzed in terms of other epistemic or mental phenomena and non-mental conditions.

### 3.3.2 Informative Partial Metaphysical Analysis of Knowledge

According to direct virtue epistemology, particular cases of knowledge and the competences which explain them are metaphysically *co-dependent*—thus the circularity of the account, and along with it rejection of the indirect virtue epistemologist’s belief-first approach of specifying the *nature* of the competences which explain knowledge in terms of metaphysically prior and independent factors (belief and further non-mental conditions such as truth, natural causation, safety, etc.).<sup>47,48</sup> However, direct virtue

---

same reason.

<sup>47</sup>There are all sorts of cases where two phenomena are mutually dependent in a non-mysterious way. For example, in order to dance a foxtrot (and many other dances), there must be a leader and a follower. To be a the leader in a foxtrot is to lead one’s follower; to be a follower in a foxtrot is to follow one’s leader. There is no way of specifying the roles of the dancers without appeal to the other, and we need not establish one rather than the other as fundamental.

<sup>48</sup>By “natural causation” I mean here causation of the sort that can occur between causal relata regardless of the agential properties (if any) of those phenomena.

epistemology does not reject the view that we can discover the facts *in virtue of which* subjects possess and manifest competences to know. Indeed, it imposes some necessary conditions on these facts in a way that explains their co-dependence and constrains correct full metaphysical analyses, thus making progress on understanding the facts in virtue of which we know.

The central idea is that, instead of there being probabilistic or other relations between belief and truth which make it the case that subjects possess the competences that explain knowledge, it is probabilistic relations between the subpersonal cognitive and environmental facts that in particular cases are constitutive of knowledge that do so. (This is so even though the facts that are constitutive of particular cases of knowledge are so only if the proficiency condition is in place.)

For example, in DOUBLE TROUBLE although the purportedly relevant probabilistic facts relating Creola's believing and the truth of her beliefs are in place, the way in which the fact that  $p$  is related to her beliefs differs from the way that is intuitively knowledge-constituting. Intuitively, the informant must be reliable and cooperative in order for Creola to know. Direct virtue epistemology explains this by claiming that Creola's competence is characterized by a proficiency condition, which requires that there be a high probability that the whole manifestation conditions will obtain (including that the informant be reliable and cooperative) conditional on the operation of the subpersonal cognitive mechanisms of Creola that are part of the manifestation conditions.

This view about which probability facts establish epistemic competences changes the course of direction for further scientific and philosophical inquiry into the facts in virtue of which we know. Instead of looking for belief-forming mechanisms that reliably produce true beliefs, we should be looking for mechanisms that reliably constitute knowledge—that are probabilistically related not just to the truth of the doxastic state in question but to getting the truth in a particular way.

Thus, while not specifying exactly which kinds of cognitive mechanisms or which kinds of environmental facts are constitutive of a particular case of knowledge, it re-orientes the analytical project in a way that more appropriately captures the relation

between competence and achievement, and so will hopefully prove more fruitful.

### 3.3.3 Full Metaphysical Analysis of Justification

Lastly, direct virtue epistemology reverses the traditional explanatory order and provides a *full* metaphysical analysis of justification in terms of knowledge and competences to know. Justified belief is a disjunctive kind, the disjuncts of which are the two kinds of exercises of competences to know: characteristic manifestations (knowledge) and degenerate exercises (Gettier cases and justified false beliefs). Degenerate exercises, though failures, are still competent performances and thus have the positive epistemic status of being a competent performance of a competence to know.

Being a justified belief, however, is not itself an exercise of competence, and justified beliefs *as such* do not have the status of being competent performances. (Rather, they have it in virtue of being knowledge or being degenerate exercises of competences to know.) Nevertheless, all justified beliefs do have a common epistemic status *qua* justified beliefs—reliability with respect to knowledge.

While this epistemic status is not itself a kind of performance normativity, it is derivative of the performance normativity of achievements. Thus direct virtue epistemology explains the common positive epistemic status of justified belief in a virtue epistemological framework, even though justified belief is not itself a kind of competent epistemic performance but is instead a disjunctive kind that is posterior to knowledge.

### 3.3.4 Moving on

Of course there are still many more questions to ask about direct virtue epistemology, and about knowledge, justification, and belief. How are competences to know individuated? What distinguishes competences to know from other personal competences, and from non-personal dispositions? How is having *some* justification related to having knowledge-level justification? These questions and many more deserve close examination, but I cannot pursue them here. Nevertheless, I hope to have motivated



direct virtue epistemology and more generally the pursuit of a substantive knowledge-first epistemology. The Gettier problem teaches us not that a substantive theory of knowledge is impossible, but that it must be of an importantly different kind than has traditionally been sought.

## Chapter 4

### A Virtue Aistheology

#### Introduction

The default philosophical view is that the fundamental perceptual phenomenon is *having a perceptual experience*. What it is to see, and more generally to perceive, is explained in terms of the representational content of such experiences, and perhaps further causal relations between perceptual experiences and the environment. This is an *experience-first* orientation in the philosophy of perception (aistheology).<sup>1</sup> I will argue that we should instead pursue a *perception-first* aistheology. By investigating what it is to perceive things as they are directly, we can better understand perception, perceptual experience, and their places in the natural world.

Others have argued for the rejection of experience-first aistheology. Disjunctivists, for example, deny that perceiving things as they are can be understood in terms of a perceptual experience common to perceiving things as they are and hallucinating.<sup>2</sup> However, the experience-first approach remains by far more fruitful and explanatory than the perception-first approach. Why is that?

Experience-firsters have a tool—representational contents. They use this tool to (i) characterize the nature of perceptual experience, (ii) explain other perceptual phenomena such as perceiving things as they are, illusions, and hallucinations, and (iii) provide a programmatic way for investigating the facts in virtue of which perceptual phenomena obtain. In contrast, perception-firsters—mainly Disjunctivists and naive

---

<sup>1</sup>Surprisingly, there is no word for the study of perception as there is for the study of knowledge (epistemology). I've made one up in order to draw parallels with virtue epistemology, and hope that it will catch on. The spelling is intended to eliminate possible confusion with "aesthetics", another term derived from the same greek word ("aisthesis"), meaning clear perception.

<sup>2</sup>E.g., McDowell (1994); Martin (2002, 2006); Brewer (2011).

realists—have been largely silent on substantive questions in the philosophy of perception or have explicitly endorsed quietist methodological claims about such issues as the nature of perception, its relation to illusion and hallucination, and how it might obtain in virtue of non-mental facts.

However, there is nothing about the perception-first program that requires it to be less explanatory, or less naturalistic, than an experience-first program. Perception-firsters take perception to be metaphysically prior to both perceptual experience and illusions and hallucinations. They can characterize the nature of perception and use this account to explain related perceptual phenomena and how all perceptual phenomena might obtain in virtue of non-mental facts.

The aim of this paper is to present such theory. According to *direct virtue aistheology*, perceiving things as they are is a matter of manifesting a competence to perceive things. Perceiving is a kind of activity in which the perceiver is directed at what she sees as the *target* of her activity.<sup>3</sup> The nature and distinctive features of perceptual activities are determined by the competences that they are manifestations of. This view about the nature of perception allows the direct virtue aistheologist to explain many facts about perceiving and related perceptual phenomena. It also generates a substantive and fruitful naturalistic research program.

I will not develop the theory in detail here. The aim of the paper is to present a view which, like the experience-first representationalist approach, provides a programmatic way of pursuing how perception and perceptual experience obtain in virtue of non-mental facts, leaving much room for in-house debate about the details. In section 1, I discuss a puzzle that has often been thought to motivate a perception-first, non-representationalist approach: Berkeley's Puzzle. I show how one main response to it—John Campbell's relationalism—fails to be an adequately explanatory alternative. In section 2, I develop the main aspects of direct virtue aistheology, showing how the theory explains what it is to perceive an object as having a property in a way that can be programmatically investigated by cognitive science.

---

<sup>3</sup>Direct virtue aistheology shares many of the features of my direct virtue epistemology (2014b).

In section 3 I explain and discuss the theory's account of the nature of the qualitative nature of perceiving. I show how direct virtue aistheology satisfactorily explains how we can perceive things *as* having properties in a way that does not appeal to attribution or representation of those properties. I also explain how the account is a nonconceptual theory of perceiving *as*, which makes clear how perceptual experience can be nonconceptual and nevertheless rationalize beliefs. In section 4 I explain how direct virtue aistheology can explain the nature of hallucinations, illusions, and perceptual experience more generally in terms of competences to perceive, and can predict that perceptual experiences in cases of illusion and hallucination may be subjectively indiscriminable from cases in which we see things as they are.

#### 4.1 Berkeley's Puzzle and Adequate Explanatory Scope

Berkeley's Puzzle seems to be generated by a very plausible starting point.<sup>4</sup> The idea is that causal relations to a perceiver's environment make her perceptually aware, or seem to be perceptually aware, of her environment. This plausible starting point is almost always clarified as the claim that in having perceptual experiences we seem to be aware of our environments because of causal relations between our *perceptual experiences* and our environments. These causal relations establish representational facts between our experiences and our environments, and we seem to see things as being a certain way because our experiences represent them to be that way.<sup>5</sup> This naturalistic, representationalist approach to explaining the directedness of perceptual experience is the foundation of the contemporary experience-first approach. It is at least as old as Locke's *Essay Concerning Human Understanding* (1690), and so is what many take to be the central challenge to it.

Berkeley (1710) argued that even if all of the causal relations were as Locke supposed, they would not suffice to explain how *perceivers* have experiences that are as

---

<sup>4</sup>While there are other experience-first approaches than representationalism (e.g. adverbialism), I shall focus in this paper on how direct virtue aistheology provides an alternative to experience-first representationalism, because it is the dominant approach.

<sup>5</sup>Dretske (1986); Fodor (1987); Millikan (1989), etc. Nearly all representationalists about perception think something like this is right.

of their environments.<sup>6</sup> Why should the fact that a certain state of affairs caused (or tends to cause) the subject to have one experience rather than another suffice to make her aware, or seem to be aware, of its cause? Why should being in a state which is related to the environment in this contingent way suffice to make *its bearer* perceptually aware of her environment? Here John Campbell and Gareth Evans each express this worry:

Here is one way to put Berkeley's Puzzle: merely appealing to the idea of experience as sensations caused in us by the world will not explain how it is that we have the idea of physical objects as the subjects of external interventions [i.e. as mind-independent objects].

—Campbell (2011), p. 37

[I]t is quite obscure how, if one mental state represents a particular in virtue of one sort of causal relation to it, and another mental state represents a particular in virtue of another sort of causal relation to it, that the sheer difference between the causal relations could generate a difference in *content* between the two mental states, given that it need not impinge on the subject's awareness.

—Evans (1982), p. 83

The central difficulty may be expressed as follows: Why does an experience's relations to the environment change the nature of the *subject's first-personal* relation to the environment? The directedness of perceptual experience seems to be distinctively mental, *originally intentional*. In perception we seem to be *aware* of things in our environment as being a certain way. But on the naturalistic story we are considering, what it is for an experience to have content has nothing to do with what it is for the subject to have the experience. It thus becomes mysterious why the experience's having the content that it does should be or constitute a distinctively mental fact about the experience.

---

<sup>6</sup>There have been several variations of this argument in recent years (Evans (1982); McDowell (1994); Campbell (2002b); Travis (2013); Robinson (1994)), but for our purposes, we may treat the argument as unified. One may even view Searle (1984)'s argument against Fodor's Computational Theory of Mind as a version of Berkeley's Puzzle.

The difficulty runs very deeply in contemporary naturalistic theories of perception. It arises for any experience-first approach that endorses a Lockean naturalistic strategy, even representationalist approaches on which perceptions and hallucinations have different contents.<sup>7,8</sup> Nor is the problem restricted to representationalist accounts on which representational relations hold in virtue of causal relations, as the quotes from Campbell and Evans above might lead one to believe. The problem is due to the Lockean way of “precisifying” the intuitive starting point we began with. The representationalist replaces the intuitive claim that we have experiences as of our environments because of causal relations we bear to our environments with the claim that we have experiences as of our environments because our *experiences* bear causal relations to our environments. This move already supposes that perceptual experiences *qua* mental states are metaphysically independent of one’s environment, and that the representational relation may be established between the experience as representans and content as representatum.

This move is, of course, intended to facilitate the naturalistic project. If we can specify the relationship between representans and representatum without appeal to the distinctively mental properties of the representans, then we have a good shot at naturalizing it.<sup>9</sup> But this is precisely the move that opens the naturalist up to Berkeley’s Puzzle, for we may ask: Suppose that the representational facts are exactly as the representationalist suggests. Why should (naturalistically respectable) representational facts between an experience and its representational content constitute, or be a matter of, *the perceiver* being directed at what her experience represents?<sup>10</sup>

Berkeley’s Puzzle has remained a central difficulty for experience-first naturalistic

---

<sup>7</sup>On these views, there is a mental representational state, or exercise of a representational capacity that is exactly the same in cases of perceiving a particular object and hallucinating. This must be so because the relation between experience and object perceived is to supposed to be *representational* in the sense that aids the naturalistic project. Thus the representing state must be metaphysically independent of its bearing representational content.

<sup>8</sup>Some reject this naturalistic approach, but then it is unclear to what extent representationalism is *explanatory*, rather than a mere terminological variant of describing the perceptual phenomenon to be explained.

<sup>9</sup>See Field (1978) for a beautiful articulation of this strategy.

<sup>10</sup>Searle’s (1980; 1984) Chinese Room Argument may be understood as making exactly this point, although he did not intend the worry to be as general as I am presenting it here.

theories of perception, and there are only two kinds of responses to it, neither of which is satisfactory.<sup>11</sup> On the one hand, one may hold onto representationalism. One can do so by either rejecting the idea that it faces any explanatory burden to respond to Berkeley's Puzzle or by supposing that, because it seems like the only plausible naturalistic view, that there must somehow be a solution.<sup>12,13</sup> On the other hand, one may reject the representationalist naturalistic approach in favor of trying to capture facts about the subject's first-personal perspective. One can do so either by embracing idealism (Berkeley (1710); Robinson (1994)), or by espousing the claim that perception is a matter of primitive openness to one's environment (Brewer (2011); Campbell (2002b); McDowell (1994)) about which little more can be said.

Put so starkly, it should be clear that neither of these responses is satisfactory. As long as it is both plausible that perceptual directedness obtains in virtue of the natural, mind-independent facts and that such directedness is a matter of the perceiver having a distinctively first-personal perspective on her environment, defending either horn at best leads to ideological entrenchment, with those who think naturalism should be defended on all costs on the one side, and with those who think the same of perspectival facts on the other.

If the naturalist is to resolve Berkeley's Puzzle, she must explain how causal and other natural relations that hold between *sub-personal*, non-mental entities and the perceiver's environment could constitute perceiving as one unified person-level phenomenon.<sup>14</sup> We cannot divide up the naturalistic project into that of discovering what it takes for a subject to have a perceptual experience and that of discovering what it takes for an experience to have the content that it does. We must explain how natural

---

<sup>11</sup>For various recent discussions of the issue, see Campbell (2002a, 2011, 2012); Cassam (2011); McDowell (1994); Pappas (2007); Brewer (2011); Travis (2013).

<sup>12</sup>This is a claim made by several representationalists who deny a straightforward naturalistic representationalism such as Fodor's Representational Theory of Mind (Fodor (1980, 1987)).

<sup>13</sup>E.g. Burge (2010), Fodor, p.c.

<sup>14</sup>I leave open exactly what the metaphysical relation is supposed to be, whether constitution, grounding, identity, or some other relation. I will use the word "constitution" here for the sake of simplicity of exposition only.

facts give rise to a phenomenon that is both intrinsically first-personal and intrinsically directed.

#### 4.1.1 Is Campbell's Relationalism an Equally Adequate Alternative?

To make clear the sort of perception-first account that is needed, I will briefly discuss John Campbell's view. Campbell has been one of the most prominent opponents of experience-first aistheology, and has used Berkeley's Puzzle to argue for its rejection.<sup>15</sup> However, the approach that he offers to replace it with both fails to be equally explanatory as experience-first views and fails to adequately answer Berkeley's Puzzle.

According to Campbell's *Relational View*, a subject's perceiving a scene is a matter of her bearing a certain kind of relation to the scene. The relation, he claims, is "primitive", "simple".<sup>16,17</sup> It cannot be further analyzed, and he does not say much more about it. What he does say can be summarized as follows. First, the relation that the perceiver bears to her environment ("the scene") causes and justifies her grasp of demonstrative concepts that refer to objects in the scene, as well as her possession of empirical concepts generally<sup>18</sup>. Second, he claims that "the qualitative character of experience is constituted by the qualitative character of the scene perceived".<sup>19</sup>

While we agree on the first point—the fact that experience causes and justifies various aspects of thought—I think that this is an *explanandum* (and so do many of Campbell's representationalist opponents). The question we want to answer is: what is the nature of the perceptual relation, and how does it obtain in virtue of non-mental facts, so that it can play this explanatory role? Taking this role to be a primitive fact about the perceptual relation is thus unsatisfactory.

Second, it is highly implausible that the qualitative character of experience is fully constituted by the qualitative character of the scene perceived, for the simple reason

---

<sup>15</sup>Campbell (2002a, 2011, 2012).

<sup>16</sup>Campbell (2002b), p. 117 and p. 114, respectively.

<sup>17</sup>See McLaughlin (2010) for further discussion.

<sup>18</sup>Campbell (2002b), e.g. p. 16, Campbell (2011), p. 47

<sup>19</sup>Campbell (2002b), p. 114.



that it becomes quite mysterious how hallucinations and illusions could be indiscriminable from cases when we see things as they are. If the qualitative character of experience is constituted by the scene perceived when we perceive things as they are, what is it constituted by in cases of hallucination and illusion? How might illusions and hallucinations be indiscriminable from cases in which we perceive things as they are, given that they are totally different kinds of perceptual phenomena?<sup>20</sup>

In contrast, the experience-first approach (at least at first pass) has a simple explanation. The qualitative character of experience is determined by its representational content, which it may have in both cases of perceiving things as they are and in cases of hallucination and illusion.<sup>21</sup> We need a perception-first theory that has the same explanatory scope as experience-first theories. It should explain how perception can cause and rationalize beliefs as well as how it is possible for cases of illusion and hallucination to be indiscriminable from cases when we perceive things as they are.

Campbell—to his credit—is one of the only perception-firsters to even address the question of how perception might obtain in virtue of natural non-mental facts. In doing so, he provides the following analogy:

But if you are caught by the idea that the existence of brain processing in vision means that a Representationalist view must be correct, it may be helpful to consider another analogy. Suppose we have a medium which, like glass, can be transparent. But suppose that, unlike glass, it is highly volatile, and needs constant adjustment and recalibration if it is to remain transparent in different contexts. ... The upshot of the adjustment, in each case, is still not the construction of a representation on the medium of the scene being viewed; ... the medium becomes transparent. You might think of visual processing as a bit like that. ... there is a kind of complex adjustment that the brain has to undergo, in each context, in order that you can be visually related to the things around you; so that you can see them, in other words.

—Campbell (2002b), p. 119

While Campbell admits that the analogy is “homuncular” (p. 119) and so is not to be taken too seriously, it is extremely difficult to understand how the analogy is supposed

---

<sup>20</sup>See Siegel (2004a, 2008) for this kind of criticism. Also see Pautz (2011) for related worries about Campbell’s relationalism.

<sup>21</sup>Those who think that perceptual experience has or can have singular content must modify this story.

to help us understand the relation between cognitive processing and perception in a non-homuncular way. According to the naturalist view as I am understanding it, non-mental subpersonal cognitive and environmental facts *constitute* perception and perceptual experience. We want to understand how this is so. Campbell's claim that there is a kind of "complex adjustment" of the brain that in the appropriate context enables perception is not a step towards a naturalistic *explanation* of perception.<sup>22</sup> What we need is an account of the perceptual relation that allows for a systematic way of relating aspects of perception and perceptual experience to natural, non-mental facts. Only an account of this sort will satisfactorily resolve Berkeley's Puzzle, because only an account of this sort will show how natural facts might give rise to a perceiver's awareness of her environment.

In the rest of the paper, I show how direct virtue aistheology can discharge these explanatory burdens. In section 2, I explain the main commitments of direct virtue aistheology and show how it provides a much more attractive and substantive account of the relation between perception and cognitive processing, without encountering the same Berkeleyan worries that experience-first accounts face. In section 3 I show how direct virtue aistheology's account of the perceptual relation allows for a plausible account of the qualitative nature of perceiving. In section 4, I develop the account of perceptual competences to extend the account of qualitative character to cases of hallucination and illusion.

## 4.2 The Perceptual Relation and Scientific Investigation

Perception, according to direct virtue aistheology, is a kind of activity, an activity essentially performed by a perceiver. What the perceiver perceives is the *target* of this activity, not what is represented by her experience. Perception is thus a species of the more general person-level event kind *activity-with-a-target*. By conceiving of the perceptual relation as an agent-target relation, as opposed to an experiencer-experience-representatum relation, the direct virtue aistheologist denies that there is any such

---

<sup>22</sup>McDowell (2010) also suggests that the perceptual capacities that are studied by cognitive science merely play an "enabling" role in perception.

thing as an experience distinct from engagement in the target-directed activity, and so denies that there is any such thing as an experience that has status *qua* mental, *qua* first-personal without bringing in facts about perceptual directedness. There is thus no room for the distinction between the perceptual mental state and what it is about that engenders Berkeleyan skepticism.

#### 4.2.1 The Perceptual Relation

The idea that perception must be agential in order to avoid Berkeley's Puzzle is not new.<sup>23</sup> However, those who have pursued this approach have employed an overly strong, overly intellectual, account of agency which itself raises problems. Berkeley's idealism, Kant's appeal to *Concepts*, and even McDowell's claim that "there is no need to look for priority [between mind and world] in either direction" all engender skeptical worries of a different sort, namely that the world we have access to through perception is not mind-independent as we suppose it to be.

Is there an account of perception as a kind of activity that does not face this problem? Perception cannot be a matter of creation or *spontaneity*; nor is it an intentional action of any kind, in the sense normally meant by philosophers. Intentional actions are often thought to be things that we can decide whether to do or otherwise control, that we can reflect upon our reasons for so doing, that are subject to "Why?" questions, and so on.<sup>24</sup> It is highly implausible that seeing satisfies these requirements. However, it is equally implausible that all activities with agents, aims, and targets satisfy these requirements. For example, when many non-human animals perform the activities of eating, mating, fleeing, hunting, etc., we think that they engage in purposeful activities with targets (food, mates, predators, prey, etc.) even though these animals do not perform intentional actions in this more robust sense.

Indeed, even in some human cases where these further features are absent, it is

---

<sup>23</sup>Berkeley (1710); McDowell (1994); Kant (1997).

<sup>24</sup>For an overview see Wilson & Shpall (2012). Anscombe (1957) introduced the applicability of "Why did you do that?" questions (that ask for reasons for action) as a criterion of intentional action.

highly plausible that humans are engaged in purposeful, targeted, activities. For example, we still count the addict as lighting up her cigarette even though she does it on impulse, reflectively rejects this action, and would rather not perform it. She still counts as the agent of her activity in the sense at issue; she still acts with the purpose of lighting up her cigarette, and the cigarette is still the target of her activity.

It is this very basic sense of activity that I mean when I claim that perceiving is a kind of activity.<sup>25</sup> Perceiving is a matter of the agent engaging with an object *qua* target of her perceiving activity. This is all that is meant or required by the claim that perception is an activity with an agent, aim, and target, just as with these other cases. Perceiving is among the activities that animals (including ourselves) just do, often whether we want to or not, and often without much control over when and how we do it.<sup>26</sup> The sort of agency required to provide an account of perceptual directedness is not *spontaneity* as Berkeley, Kant, and McDowell supposed, but rather, it is just engaging with mind-independent objects as targets of one's activities.

If perception is an activity with an aim and a target, then the naturalistic project becomes one of explaining how non-mental facts give rise to activities of this sort, in a way that explains the distinctive features of perception and perceptual experience. If there are tools that the perception-firster can use to facilitate and make plausible the carrying out of this project, then the naturalist really can satisfactorily answer Berkeley's Puzzle. According to direct virtue aistheology, competences are exactly the tools we need.

#### 4.2.2 Scientific Investigation of Competent Performance

Direct virtue aistheology claims that perceiving things as they are is a manifestation of a competence to perceive. In this subsection, I show how the view fits neatly into

---

<sup>25</sup>In all of these cases, "I didn't know I was doing it" is not an acceptable response. However, because "I couldn't help it" is such an obvious rejoinder, it may seem strange to ask the question. Anscombe's famous test of intentional actions, it seems, picks up more cases than she was looking for. As long as "I just was" or "I couldn't help it" are ways of asking the question, "Why?" questions delineate the realm of purposeful activities, not just intentional actions.

<sup>26</sup>There is, however, a gray area here. We can decide where to look and what to think about, and we can have reasons for doing so.

contemporary empirical investigation of the mental. The main difference between the direct virtue aistheologist's approach and the experience-first approach is as follows. Causal and environmental regularities, instead of constituting representational facts, constitute facts about competence possession and exercise. The mechanisms that cognitive scientists investigate are *bases* of competences. On a particular occasion the operation of the basis of the competence, together with the obtaining of certain environmental facts, constitutes the exercise of competence, i.e. perceiving things as they are.

Some theorists argue (independently of the considerations about perception discussed here) that computational explanations are already explanations of person-level competent performance.<sup>27</sup> For example, Frances Egan (2013) argues that computational specifications of cognitive mechanisms accomplish two distinct explanatory tasks. First, in providing a non-relational (formal or mathematical) specification of a cognitive mechanism, one explains how the mechanism works *qua* mechanism. Since we know how to build mechanisms that compute functions and manipulate symbols, formally or mathematically specifying a mechanism helps us to understand the physical nature of the mechanism. Second is the task of showing how the operation of a mechanism so specified makes it the case that the cognitive task—specified in common-sense pre-theoretic terms—is reliably accomplished.<sup>28</sup> She argues that attribution of worldly content is guided largely by pragmatic factors, and whose purpose is to help *us* understand how the mechanism's computing the mathematical function, together with the right environmental circumstances, constitutes the competent performance we are trying to explain:

[T]he contents are assigned ... as a way of ... helping us keep track of changes in the system caused by both environmental events and internal

---

<sup>27</sup>See also Chomsky (1995, 2003) but beware. His competence/performance distinction *does not* map onto my competence/manifestation distinction (See Chomsky (1965).) A full account of a perceptual competence in my sense will also explain intricacies of performances, both manifestations and distinctive failures.

<sup>28</sup>She argues along similar lines in earlier work as well. See esp. Egan (1995). She argues for the stronger thesis that all such explanations are mathematical explanations. I do not need to commit myself to this claim here. It may be sufficient, e.g., that one have a translation procedure from 1s and 0s to the language of a computer program.

processes, all the while with an eye on the cognitive capacity (e.g. seeing what is where) that is the explanatory target of the theory.

—Egan (2013), p. 11

If it is true that content is attributed to computational symbols for the purpose of explaining how computationally specified mechanisms underlie the competent performance in question, then interestingly computational theories of cognitive performance are actually already situated in a *competence framework* for thinking about mentality. We pre-theoretically identify what it is that the person does—e.g. seeing or reaching for and grasping objects—and try to provide a computational explanation of internal processes that underlie such performances, so that coupled with facts about the environment and causal regularities we can understand how people do what they do.

Although Egan is skeptical that a non-pragmatic account of the content of computational symbols can be given (and I am sympathetic), it is worth noting that direct virtue aistheology need not commit to this view. As long as the explanatory purpose of attributing content to symbols is to explain competent performance as Egan argues, the direct virtue aistheologist is in a good position to adopt current practice of computational cognitive science pretty much as it is.

It is important to point out, however, that direct virtue aistheology does not *commit* to any account of cognitive processing. It is compatible with, e.g., a dynamical systems account as well as a computational account.<sup>29</sup> Nevertheless, I use Egan's account of explanation in computational cognitive science in order to make it clear that there is no reason why a perception-firster, or anyone who denies that having a perceptual experience *is* a matter of being in a state with a certain representational content, need deny that perception *is constituted* in part by the operations of sub-personal computational cognitive mechanisms, which include symbols with representational content. What they must deny is that somewhere down the line there is a state with representational content that is the output of cognitive processing, and which just is the *person-level* mental perceptual experience. Egan's account is an attractive alternative

---

<sup>29</sup>van Gelder (1998).

way of understanding the relation, one that fits nicely with direct virtue aistheology.

There are many independent reasons to endorse Egan's account. Here I briefly discuss one. It may turn out that there is no output of visual processing that is a good candidate for being the perceptual experience (i.e. that has the representational content that appropriately corresponds to what we visually experience.)<sup>30</sup> However, if perception is a competent performance, and the cognitive processing merely partially constitutes such performances, then it is not required that certain computational symbols represent either the objects that we see or the properties that we see objects as having. All that is required is, to use an example of Dennett's, that we may fully understand why the computer program thinks it should get its queen out early in terms of its underlying computational processing and relations to the environment, even though nowhere is there any state of it that represents that it should get its queen out early.<sup>31</sup> Cognitive processing may explain person-level phenomena in a more complex way.

In particular, features of the way in which computational operations *treat* symbols may have an integral role to play in the explanation of person-level directedness. Here is an example from the current psychological literature on visual attention. There is a wealth of evidence that at least some visual attention is object-directed. Among other results, test subjects are faster at recognizing whether two properties are instantiated in a display if they are instantiated in the same object as opposed to different objects, even controlling for location.<sup>32</sup> There is evidence that attention spreads over objects,<sup>33</sup> even those that are partially occluded.<sup>34</sup> Visual short term memory seems to "chunk"

---

<sup>30</sup>Dale Purves and his lab have proposed one of the best current scientific explanations of why we perceive the relevant lines as being of different lengths. This theory does *not* involve a computational process resulting in the attribution of different lengths (or of relative difference in length) to the lines. Rather, it supposes that cognitive systems compute highly abstract statistical information about retinal stimuli and features of our phenomenal awareness correspond to these statistical properties. (See Howe & Purves (2005b,a); Purves et al. (2011, 2001); Purves & Lotto (2002).) I discuss this example in more detail in Miracchi (2014a).

<sup>31</sup>Dennett (1978), p. 107. Fodor (1987) discusses this passage on pp. 21-3.

<sup>32</sup>Baylis and Driver (1993), building on work by Duncan (1984). Blaser et al. (2000b), 2000a obtained similar results using superimposed Gabor patches. See also Farah (1989) for similar evidence.

<sup>33</sup>Rock & Gutman (1981).

<sup>34</sup>Moore et al. (1998).

information in a way that is object-based, rather than feature-based.<sup>35</sup> Several studies have revealed that some priming is object-based, not location-based.<sup>36</sup> Data on visual neglect patients also suggest that visual neglect is object-based.<sup>37</sup>

Interestingly, computational explanations of object-directed visual attention involve computational operations over both representations of objects and representations of properties. What explains why visual attention is object-directed (as opposed to object-and-property-directed) is not only a matter of the contents of certain representations it computes, but also the way computational operations treat the various representations they compute over and how that treatment of representations relates to environmental situations.

And yet there is no mystery here. We are well on our way to understanding how the activity of visually attending to objects has all of the features that it does in terms of the operations of sub-personal computational mechanisms and their relations to the subject's environment. It is exactly this sort of computational explanation of a directed activity that I am claiming perceiving objects (often thought to be prior to visual attention) deserves.

The preceding discussion supports the view that we can conceive of perception as an exercise of competence, not identify it with a representational state of the cognitive system, and yet still investigate how cognitive processing and environmental conditions give rise to perception in a rigorous scientific way. Direct virtue aistheology, then, adequately answers Berkeley's Puzzle. It provides an account of the perceptual relation that does not engender worries about how the perceiver might stand in such a relation and nevertheless fail to be aware of the mind-independent world, and it

---

<sup>35</sup>Luck & Vogel (1997), Jiang & Chun (2001).

<sup>36</sup>Kahneman et al. (1992) presented subjects with two square figures with letters in them. Then the letters disappeared and the boxes moved. One of the letters then reappeared, either in the same box as it was originally in, or in the other box. Subjects were faster at recognizing the letter when it appeared in the same box. Tipper et al. (1991) showed that this phenomenon occurs even when the objects switch locations, showing that the priming cannot be location-based.

<sup>37</sup>Behrmann & Tipper (1994), and (1999), and Tipper & Behrmann (1996) showed that the neglect of left-neglecting patients is object-based, rather than merely egocentric-location based. If something is neglected, the means that the subject is not able to attend to it. If it is objects that are neglected, and not mere visual regions, then that suggests that it is objects that are attended to.



also makes it plausible that natural natural facts establish this perceptual relation, connecting up with scientific investigation of perception. What is left is to determine whether direct virtue aistheology can provide an account of perceptual experience that can match experience-first theories in explanatory scope.

### 4.3 The Rational Role of Perception

When a subject perceives an object, she perceives it in a certain way, as having certain properties. How can the perception-firster provide an account of what it is to see *o as F* in a way that explains how perception rationalizes belief (and action) in the distinctive way it does? Instead of thinking about ways of perceiving as things that perceivers grasp (i.e. contents), I suggest that we think about them as things that perceivers *instantiate*. What it is to see *o as F* is not to grasp or represent a way of perceiving *o*; it is to engage with *o* in a certain way, a way which is fully determined by the nature of the competence manifested. How might this be so? How, in other words, can we explain how seeing *o as F* is just a way of seeing *o*, not a matter of attributing *F* to *o*?

There is a quote of Gareth Evans' that Jason Stanley has recently appealed to in support of the role of (Fregean) contents in explaining understanding, and it will be a useful starting point for contrast:<sup>38</sup>

A way of thinking of an object is no more obliged to get in the way of thinking of an object, or to render thinking of an object indirect, than is a way of dancing liable to get in the way of dancing, or to render dancing somehow indirect.

—Evans (1985), pp. 302-303

I think that Evans is right, as long as this is understood as the mere rejection of a necessity claim. But it does not support a Fregean account of ways of thinking, or ways of seeing. On its most natural construal, when one dances, one's relation to a way of dancing is the *instantiation relation*, so of course it doesn't "get in the way". We

---

<sup>38</sup>Stanley, p.c. See also Stanley (2011) for discussion of Fregean ways of thinking. Stanley's discussion bears several connections to direct virtue aistheology, but I do not have the space to address them here.

can describe and compare ways of dancing, but we should not think that the nature of dancing is being related to a certain kind of abstract object or that the relevant abstract object is doing real work in explaining the directedness of actions and the kind of action being performed. What does the explanatory work is the nature of the activity and the facts that make it the case that one is engaged in the sort of activity one is engaged in.

I think that the same is true for perception. Just as dancing is an activity with an aim, so is perceiving. And just as there are various ways of performing the activity of dancing, so there are various ways of performing the activity of perceiving. Ways of perceiving may themselves be abstract objects, but the relations we bear to them are instantiation relations, not *grasping* relations. What it is to perceive in a certain way is not to grasp or represent a way of perceiving; it is to engage with what one perceives in a certain way. While appeal to representational vehicles and contents are useful tools for modeling and systematizing perceptual activities, and should retain their status in theories as such, according to direct virtue aistheology they do not illuminate the nature of perceiving or its place in the natural world.<sup>39</sup>

How, then, might we think about ways of perceiving objects, in a way that explains why we perceive objects as having properties? Instead of appealing to demonstrative concepts as a guide to the nature of perception, as several theorists do, we can appeal to indexical concepts, in particular, some suggestions made by David Kaplan.<sup>40</sup> According to Kaplan, when indexical concepts are exercised, they just pick out objects. So, e.g., in thinking a *here* thought, one just thinks about the particular place where one actually is. However, Kaplan thought that the fact that indexical concepts pick out particular objects in a way that systematically depends on the context of thinking could make a difference directly to the *way* in which one thinks about the object,

---

<sup>39</sup>I am thus not denying that perceptual experiences have accuracy conditions in *any* sense. (Susanna Siegel (2010) does an excellent job in arguing that this view is untenable.) I am denying that appeal to accuracy conditions can help us to explain the nature of perception or its distinctive features.

<sup>40</sup>While mainly concerned with indexical expressions such as *here* and *now*, he was also clearly concerned with the kind of thoughts such expressions paradigmatically express. See esp. Kaplan (1989) pp. 774-5. For examples of the use of demonstrative thought as helpful for understanding perception, see McDowell (1994); Burge (2010); Martin (2002).

and can explain why it is rational to believe and do certain things on the basis of so thinking. So, e.g., when one thinks of a place as *here*, one thinks about it *as the place where one is*—not because one attributes *being the place where one is* to the place one is thinking of, but because those are the features of the object in virtue of which one is thinking about it at all.

While Kaplan's formal treatment of indexicals has been widely influential, this idea about what explains rational relations among mental states has been largely ignored. This is not surprising given the prevalence of representationalism. If one holds that the rational contribution of a mental state is exhausted by its representational content, it is difficult to see how metasemantic facts could make a difference to one's mental life without making a difference to what content one entertains.

However, direct virtue aistheology can incorporate this idea in a straightforward and intuitive way. Clearly the nature of a competence can make a direct difference to the features of its manifestation. Indeed, it is quite plausible that the competence manifested fully determines the features of its manifestation. We may suppose then, that perceptual competences are competences to see things of certain kinds, and that they issue in manifestations that are ways of perceiving things of those kinds *as being of those kinds*. Just as (to use competence terminology) the manifestation of a competence to think thoughts about the place where one is issues in manifestations that are ways of thinking about things *as the place where one is*, so are perceptual competences. That is, what it is to see *o* as *F* is to manifest a competence to see *F* things.<sup>41,42</sup>

Perception neither involves the attribution of properties to objects nor requires that perceivers have higher-order abilities to attribute properties to objects seen. Instead,

---

<sup>41</sup>That is, the full good case of seeing *o* as *F* is manifesting a competence to see *F* things. Veridical illusions are possible, and I will explain how they may be dealt with below.

<sup>42</sup>The account also leaves open the possibility that, e.g., pure phenomenal redness is experienced solely in virtue of "in the head" non-mental facts. There is no reason why subpersonal cognitive features in virtue of which the subject possesses and manifests a competence to perceive may not also make a difference to the qualitative character of her perception. What is important for the direct virtue aistheologist is that perceiving is not constituted piecemeal, i.e., by appeal to such conscious or other mental states and relations that those states bear to the subject's environment. The facts in virtue of which one has an experience of, e.g., phenomenal redness must be related to other sub-personal facts so as to generate a unified perceptual activity. Thus the possibility of inverted color spectra (Block (1990)) do not present a difficulty.

the properties of an object that are exploited by the competence exercised in perceiving it. These properties make a difference to the *way* in which one perceives it because they partially determine the nature of the competence exercised.

Direct virtue aistheology thus explains seeing *o* as *F* in entirely perception-first terms. It is the fact that perceptual competences are competences to perceive objects of certain kinds that explains why a manifestation of that competence is a case of perceiving an object *as being* of that kind. There is no need to appeal to representation, attribution of properties, or content of any kind, beyond appealing to the object seen as the target of the perceptual activity.

This is not just advantageous for perception-firsters. It is also advantageous for those who have been arguing that perception must be non-conceptual. It has seemed to many quite implausible that our perceptual capacities are as sophisticated as our thought capacities—that perceptual experiences themselves involve the application of concepts to what we see. A standard (and plausible) way of understanding this claim is that perception does not involve *predication* of the kind thought involves. The capacity for predication seems to be a kind of capacity that is too sophisticated to be operative in perception. At a minimum, predicating *F*-ness of an object seems to require that one be able to independently identify the object in question, but in perception it seems that what properties one sees an objects as having help you to see the object at all.<sup>43,44</sup>

However, it is quite difficult to articulate what nonconceptual perception might be like, if one has an experience-first, representationalist approach. On that approach what it is like to perceive an object is fully determined by the content of the perception, and so one seems forced into thinking of *seeing as* as involving some kind of attribution or predication. What else could seeing *a* as *F* be except the entertaining of some content to the effect that *a is F*, and how can this not involve predicating *F*-ness of *a*?

---

<sup>43</sup>Burge (2010) argues for this in detail.

<sup>44</sup>Depending on who you talk to, predication requires a lot more than this. E.g., Evans (1982); Sellars (1956); McDowell (1994).

Tyler Burge tries to motivate the difference by supposing that in perceptual experience attribution is of the form *that F* as opposed to *that is F*. But it is not clear how to interpret this idea, for when we usually say “that green cup”, we give our interlocutors a description that helps them identify what we want to talk about, but which they could already identify prior to the description. Burge merely claims that there is a kind of nonconceptual, identificatory, attribution, but does not give any support for this claim.<sup>45</sup>

In contrast, on direct virtue aistheology it is quite clear how the properties one sees a thing as having play an identificatory role. Indeed, that is the whole point. The properties that one sees a thing as having are the properties that, in the case where one manifests a competence, make it the case that one is and continues to perceive the object.

Moreover, by using indexicals as a model, direct virtue aistheology can adequately respond to the traditional worry for nonconceptualists about perception—namely that only if perceptual experience were conceptual could it rationalize beliefs.<sup>46</sup> Let us briefly return to the analogy with indexical thoughts. Consider the following inference:

### **Singular**

- (1) Here is a beautiful statue.
- (2) A place exists.

While the inference is a little bit odd, it is rational *from the subject's perspective*, and it is so in virtue of the way in which the subject thinks about the place she is referring to as *here*. That is, for subjects that in fact do have the concept of *being a place*, the kind of sensitivity one has in thinking about an object as *here* rationalizes the belief, of the object thought about, that it is a place.

---

<sup>45</sup>Nor does current semantic theory weigh in Burge's favor. The majority of contemporary theories of demonstrative utterances hold that the “F” in “that F” plays a predicative role if it constrains the possible referent. (See Braun (2007) for an overview.)

<sup>46</sup>E.g. McDowell (1994), Ginsborg (2009), Cassam (2011), Bengson et al. (2011), Brewer (1999), Brewer (2005).

Note also that the explanation of why it is rational to infer the existential claim that a place exists does not rely at all on the fact that “is a beautiful statue” is predicated in (1). What explains the rationality of the move from a *here* thought to an existential claim about some place are facts about how one thinks about the place one thinks of as *here*. Nothing about this explanation requires the possibility of predicating properties of what one thinks of as *here*. If this is true in the case of indexicals, then there is no reason why it should not be true in the case of perception.

The analogy with indexical thoughts allays the worry that direct virtue aistheology’s account of what it is to perceive an object as having a property is not genuine *perceiving as*.<sup>47</sup> Since it is agreed that indexical thoughts about the self, here, and now rationalize beliefs about the persons, places, and times for those with such concepts, there is no good reason to think perception may not similarly rationalize beliefs. Thus direct virtue aistheology can adequately capture the sense in which we see objects as having certain properties and how doing so may nonconceptually rationalize beliefs.

#### 4.4 Hallucinations, Illusions, and Subjective Indiscriminability

In this section, I will show how direct virtue aistheology can account for hallucinations and illusions in a way which explains how they may be subjectively indiscriminable from cases in which we perceive things as they are. In order to do so, I must make some more commitments about the nature of perceptual competences and their exercises.

##### 4.4.1 Perceptual Competences

A perception-first virtue aistheology requires a perception-first account of perceptual competences.<sup>48</sup> The account I endorse differs from the two main existing approaches, in the way one might expect.

---

<sup>47</sup>I think this also provides us with the beginnings of an account of concept acquisition, although this is not the space to pursue this approach. Here I only point to the fact that our indexical thoughts have the kind of rationalizing roles where perspectival sensitivity to certain properties in thinking about an object rationalize attributing those properties to the object.

<sup>48</sup>What follows is the application of my *dual exercise account* (Miracchi (2014a)) to the case of knowledge.

The dominant view of perception that invokes competences (abilities/ capacities) is an experience-first view on which a subject perceives in virtue of exercising perceptual competences to veridically perceptually represent. Exercises of competences are cases of perceptually representing (which are taken to be or constitute perceptual experiencing). Perception is then explained in terms of an exercise of such a competence, veridical representation, and perhaps further causal relations.<sup>49</sup> Illusions and hallucinations are cases where the competence is exercised—and so the subject has an experience with a representational content—but for some reason either the content is false or the experience fails to bear the right causal relations to the subject's environment.

This account is powerful for many reasons, not least of which because it allows for a plausible account of the indiscriminability of cases of illusion and hallucination. If perceptions and hallucinations involve the same exercise of perceptual competences, and the subjective character of experience is determined by the exercise of one's perceptual competences, then clearly hallucinations and illusions can be indiscriminable from perceiving things as they are because their subjective character can be determined by the same exercise of competence.

On the other hand, there are those who reject this experience-first approach and instead claim that perceptual competences (capacities/abilities) are competences to perceive things as they are.<sup>50</sup> This approach, however, remains silent on the nature of illusions and hallucinations, and their relation to perceiving things as they are. Instead, we need an account of perceptual competences on which they have some exercises that are cases of perceiving things as they are, but are nevertheless *fallible*; that is, they can also have exercises that are constitutively failures to perceive things as they are.

The account I will now present can accomplish both of these tasks. It is an account of competences to perceive things of certain kinds, on which exercises of such competences are constitutively cases of perceiving things of those kinds and others are

---

<sup>49</sup>Burge (2011); Schellenberg (2013).

<sup>50</sup>Millar (2008); Williamson (2000).

constitutively failures to perceive things of those kinds. All we need are three fairly simple necessary conditions.

The basic idea is this: In cases where a perceptual competence is manifested, there are sub-personal cognitive mechanisms that are operative in certain environmental conditions. We may suppose that a necessary condition of the perceiver possessing a competence to perceive things of a certain kind is that there will be a certain probabilistic relation between the full conditions that in a particular case are constitutive of perceiving a thing of that kind, and the operation of cognitive mechanisms that is its proper part. That is, competences to perceive are, at least in part, established by probabilistic relations among the sub-personal non-mental and environmental facts, not by probabilistic relations among a mental state or exercise of competence and some kind of success condition (e.g. veridicality). We can state the view more precisely as follows.

Let us suppose that we are placing conditions on a competence  $C_{PF}$  to perceive  $F$  things.<sup>51</sup> We may associate with such a competence  $C_{PF}$  a way of perceiving  $W_{CPF}$ , which is its characteristic manifestation.<sup>52</sup> For example, the competence to see bounded, rigid, moving objects has seeing a bounded, rigid, moving object (*as such*) as its characteristic manifestation. We may then specify the *manifestation conditions* of  $C_{PF}$ , which are the conditions that in a particular case (against a background of possession of  $C_{PF}$ ) constitute seeing an  $F$  in way  $W_{CPF}$ :<sup>53</sup>

**Manifestation Conditions:** The manifestation conditions of  $C_{PF}$  are whatever operations of subpersonal cognitive mechanisms and external conditions together (against a background of possession of  $C_{PF}$ ) constitute a particular case of perceiving an  $F$  in the way characteristic of the competence ( $W_{CPF}$ ).

---

<sup>51</sup>The account I invoke here is an application of a more general account of competences I develop in Miracchi (2014a).

<sup>52</sup>Recall section 3.

<sup>53</sup>Because perceiving is by nature a manifestation of competence, manifestation conditions only are such against a background of competence possession.



A competence to perceive is **manifested** just in case its manifestation conditions obtain.

Clearly, this condition is not very informative. Making it more informative is a good project for further study, but my purposes here may be accomplished without doing so. The point of specifying the manifestation conditions is not to illuminate *them*, but to explain how there may be competences *to perceive things*, that have perceiving things as their characteristic manifestations, but that also can have exercises that are constitutively failures.

Of the manifestation conditions, we may isolate just the subpersonal cognitive mechanisms whose operations in a particular case partially constitute the subject's perceiving an  $F$  in way  $W_{CPF}$ :

**Basis Condition:** The basis of  $C_{PF}$  is fully constituted by the subpersonal cognitive mechanisms of the subject  $S$  whose operations partially constitute  $S$ 's perceiving an  $F$  in way  $W_{CPF}$ .

For my purposes here, I may also remain non-committal about the nature of these mechanisms and their operations (recall section 2). Just by distinguishing the basis of a competence to perceive from its manifestation conditions, we now have the resources to impose a different kind of reliability condition on perceptual competences than is normally imposed:

**Proficiency Condition:** The proficiency condition of  $C_{PF}$  requires that the objective probability of the manifestation conditions obtaining conditional on the basis of the competence being operative be sufficiently high. I.e.,  $Pr(M|O_B) \geq n$ , for some sufficiently high  $n \in (0, 1]$ .<sup>54,55</sup>

---

<sup>54</sup>For ease of exposition, I am representing the threshold as a real number  $n \in (0, 1]$ , but I need not make any commitments about how close to 1  $n$  must be, whether the threshold for the epistemic domain is the same as for other domains (e.g. baseball), or whether the threshold is determinate. (I think it is probably indeterminate.) I also need not make any commitments about whether the threshold is context-sensitive.

<sup>55</sup>While I need not commit to a particular view of probability here, I am supposing that appeal to objective conditional probabilities is less contentious than appeal to objective unconditional probabilities (see Hájek (2007) for defense of this claim), and I am supposing that some kind of non-frequentist realist

Putting these three conditions together, we arrive at the following account of perceptual competences. The proficiency condition relates the manifestation conditions to its basis, so that instead of a perceptual competence being characterized by a relation between (e.g.) a visual representation/experience and its truth/accuracy, it is characterized by a relation between the conditions constitutive of perceiving and their subpersonal “internal” proper part. Thus perceptual competences are competences to perceive, not to do anything else, such as veridically represent.

Moreover, imposing these three necessary conditions allows the perception-firster to make it plausible that how *what it is* for a subject to perceive things of a certain kind *just is* for her to manifest a competence to see things of that kind. This is because only if the manifestation conditions are related to their subpersonal internal proper part in the way specified by the proficiency condition do manifestation conditions constitute cases of seeing. (That was the reason for the caveat about background conditions in the specification of the manifestation conditions above.) This entails that perceptions are, by their very nature, exercises of competence.

Now I will show how the theory of perceptual competences can allow for degenerate exercises, and so can explain hallucinations and illusions in terms of perceiving.

#### 4.4.2 Hallucinations

The account of perceptual competences outlined above allows that they may be *fallible*. The proficiency condition also allows for another kind of exercise, one that is constitutively a failure. Because the proficiency condition does not require that the probability of the manifestation conditions obtaining conditional on the basis being operative be 1, the basis of a competence to perceive may be operative even if the full manifestation conditions fail to obtain. In such cases, we may suppose that such competences are *degenerately exercised*:

A competence to perceive is **degenerately exercised** just in case its basis

---

account of them is in the offing. I am also supposing that such probabilities are true at particular times. Changes in, e.g., causal regularities over time might result in either the acquisition or loss of competences to perceive by changing whether or not the proficiency condition is met.

is operative, but not all of the manifestation conditions obtain.

Perceptual competences, then may have exercises that are constitutively failures to perceive things in addition to exercises that are cases of perceiving. In a case where a subject's perceptual competences are merely degenerately exercised, the subject fails to perceive any object, and so hallucinates. We do not need a view on which perceiving is a matter of the object causing the subject to be in a certain mental state in order to show how the absence of an appropriate causal relation between subject and object results in hallucination.

Exercising a perceptual competence on this approach is a disjunctive kind:

A competence is to perceive is **exercised** just in case it is either manifested or degenerately exercised.

Even though what it is to exercise a competence is disjunctive, recall that it is facts about the perceptual competence exercised, i.e. stable features of the competence, that determine what one sees objects as, not actual properties exploited on a particular occasion. That is, the qualitative character is fully determined by stable features of the competence. This means that the account predicts that two manifestations of the same perceptual competence that have distinct objects as targets on the two occasions will fail to be subjectively discriminable from one another. This is so even though the states are importantly cognitively different, with the one object as target in one case and the other as target in the other. This difference does not provide the subject with any features to distinguish the two cases.

The account leaves open, then, the idea that when one degenerately exercises a competence, the activity (though a failure) has the same qualitative character as a manifestation of the same competence, and so is indiscriminable from manifestations of the same competence. In cases of hallucination, the subject degenerately exercises her competence, and facts about the nature of the competence exercised and the particular environmental conditions which obtain determine that it is a failure to perceive things in the way characteristic of the competence. Since the qualitative character of

perceiving things as they are does not depend on features of the particular object perceived, but rather on stable features of the competence manifested, there is no reason to think that the qualitative character of degenerate exercises depends on the fact that no particular object is perceived. The direct virtue aistheologist may then freely suppose that they do have the same qualitative character.

The account thus can explain the possibility of indiscernibility of hallucinations from cases of seeing things as they are in terms of sameness of qualitative character without appeal to representation, and while respecting the idea that perception is genuinely object-involving. The possibility of perceptions and hallucinations having the same qualitative character is explained by the possibility that a competence to perceive may have both manifestations and degenerate exercises. Nevertheless, cases of manifestation are cases of engaging with an object as the target of one's perceptual activity, and so is a matter of the agent making contact with the world in a way she fails to do in cases of hallucination.

#### 4.4.3 Illusions

What about illusions, then? How can direct virtue aistheology explain cases where we perceive things, but not as they are? When one perceives things as they are, one perceives objects as having many properties that they in fact have. Importantly, the facts in virtue of which one perceives objects as having these properties are not all on a par with one another. E.g., one might see something as red and see something as crimson, but the facts that make it the case that one sees it as red depend on the facts that make it the case that one sees it as crimson.<sup>56</sup> The facts that make it the case that one sees an object as a dog might likewise depend on the facts that make it the case that it seems to have a certain shape. In general, some properties of a thing may make a difference to one's perceptual engagement with it but only in virtue of others of its properties making a difference to one's perceptual engagement with it.<sup>57</sup>

---

<sup>56</sup>Note: the dependency of competences to see on other competences to see may or may not track the metaphysical dependence relations among the properties exploited by the competences.

<sup>57</sup>I am remaining neutral on what properties we perceive objects as having. If you do not think these we see objects as having such properties, please substitute your preferred example.

There thus arises the possibility that certain more fundamental aspects of one's perceptual engagement might secure an object as the target of one's activity, even while the manifestation conditions of more sophisticated perceptual competences which depend on this more basic engagement for their manifestation fail to obtain, and so are degenerately exercised. In such a case, the target of the perceiving is secured, even though more sophisticated competences are not manifested. These are cases of perceptual illusion.

Here is an example. Suppose that what seems to one to be a bird of prey is actually just a toy airplane in the distance. In this case, one has a perceptual illusion as of a bird of prey. One sees the airplane, but does not see it as it is. Still, one does see it as having some of the properties it does in fact have—it is a bounded, continuously moving object, and one sees it as such. This is enough to secure it as the target of an exercise of a perceptual competence and so to make it the case that one sees the thing, even though one sees it as a bird of prey when it is not. This is a case in which some of the perceptual competences are manifested, but higher-level ones that depend on them are degenerately exercised. The explanation of indiscriminability is exactly the same as for cases of perceiving things as they are and hallucinating. Illusions are hybrid cases of manifestation and degenerate exercise of perceptual competences, and so have their qualitative character determined by the competences they are exercises of.

#### 4.4.4 Perceptual Experience

What of perceptual experience in general, then? Cases of perceiving things as they are, hallucination, and illusion are all kinds of exercise of competence. This suggests the general thesis that to have a perceptual experience is to exercise one's perceptual competence(s). It is a disjunctive kind, because *exercise of competence* according to direct virtue aistheology is a disjunctive kind. However, as discussed above, the manifestation and the degenerate exercise of a competence have the same qualitative character. Thus all exercises of the same perceptual competence, regardless of whether they are manifestations or degenerate exercises, have the same qualitative character.

This is a clear sense in which one is in the same perceptual state in both cases of perceiving things as they are and cases of hallucination and illusion. However, having a perceptual experience is metaphysically dependent on perceiving things as they are, not the other way around.<sup>58</sup>

This account of perceptual experience also allows for a straightforward explanation of why hallucinations and illusions rationalize beliefs. Because all perceptual experiences that are exercises of the same perceptual competence have the same qualitative character, existential beliefs—such as that there is an F before one, that there is an F next to a G, and so on—are rational. (This is of course merely a subset of the beliefs that are rationalized by perceiving things as they are, which include that *that* is an F, that *that* F is next to *that* G, and so on).

Moreover, having a perceptual experience as of an F justifies the belief that there is an F before one in an externalist sense as well. This is because if to have a perceptual experience as of an F is to exercise one's competence to perceive an F thing. The proficiency condition on perceptual competence possession entails that it is highly likely that if one exercises one's competence to perceive an F thing, one manifests it. But this entails that it is highly likely that if one has a perceptual experience as of an F, then there really is an F before one (and one really perceives it). Thus even in cases of illusion and hallucination, many perceptual beliefs are justified.<sup>59</sup>

#### 4.5 Conclusion

When we see, we see things *as* being a certain way. Traditionally, the approach to providing a naturalistic explanation of this phenomenon appeals to experiential or mental states that, in virtue of bearing certain non-mental relations to the environment, represent things to be a certain way. I have here provided an alternative of account of the perceptual relation, and argued that it has the same explanatory scope

---

<sup>58</sup>Thus direct virtue aistheology is disjunctivist only on certain ways of defining disjunctivism.

<sup>59</sup>In order for perceptual experience to justify belief according to direct virtue epistemology, the competence to perceive that is exercised must partially constitute an exercise of a competence to know. I reserve discussion of this for a future date, but hopefully the reader can see how this is intended to go.

as experience-first representationalism, while avoiding some difficulties that account faces, especially Berkeley's Puzzle.

Direct virtue aistheology is a programmatic perception-first theory that can be used to investigate the features of perception, both scientifically and philosophically. It explains why in perception we see objects as having properties, and how perception has the rational role that it does. Moreover, it provides a theory of illusion and hallucination that explains their indiscriminability from cases of perceiving things as they are. Lastly, it provides an explanation of perceptual experience in general.

While there is still much work to do in developing direct virtue aistheology, I hope to have convinced the reader that a perception-first theory can be just as naturalistically plausible and explanatory as its experience-first rivals.

## Bibliography

- Anscombe, G. E. M. (1957). *Intention*. Cambridge, MA: Harvard University Press/Harvard University Press.
- Behrmann, M. & Tipper, S. (1994). Object-based visual attention: Evidence from unilateral neglect. In C. Umiltà & M. Moscovitch (Eds.), *Attention and Performance 15: Conscious and Unconscious Processing and Cognitive Functioning*, volume 19 (pp. 351–375). MIT Press.
- Behrmann, M. & Tipper, S. (1999). Attention accesses multiple reference frames: Evidence from visual neglect. *Journal of Experimental Psychology: Human Perception and Performance*, 25(1), 83–101.
- Bengson, J., Grube, E., & Korman, D. Z. (2011). A new framework for conceptualism. *Nous*, 45(1), 167–178.
- Bergmann, M. (2008). Reidian externalism. In V. Hendricks & D. Pritchard (Eds.), *New Waves in Epistemology*. Palgrave.
- Berkeley, G. (1996/1710). *Principles of Human Knowledge and Three Dialogues (Oxford World's Classics)*, chapter Treatise Concerning the Principles of Human Knowledge, Part I. Oxford: Oxford University Press.
- Bird, A. (2007). Justified judging. *Philosophy and Phenomenological Research*, 74(1), 81–110.
- Black, C. (1971). Knowledge without belief. *Analysis*, 31(5), 152–158.
- Blaser, E., Pylyshyn, Z. W., & Holcombe, A. (2000a). Tracking an object through feature-space. *Nature*, 408, 196–99.
- Blaser, E., Pylyshyn, Z. W., & Holcombe, A. (2000b). With two objects at one location, attention is “object-based”. *Presented at the 2000 meeting for the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, Florida*.
- Block, N. (1990). Inverted earth. *Philosophical Perspectives 4: Action Theory and Philosophy of Mind*, 51–79.
- Braun, D. (2007). Indexicals. *Stanford Encyclopedia of Philosophy*.
- Brewer, B. (1999).
- Brewer, B. (2005). Do sense experiential states have conceptual content? In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology*, Debate with Alex Byrne. Oxford: Blackwell, 2005.
- Brewer, B. (2011). *Perception and Objects*. Oxford University Press.



- Buckwalter, W., Rose, D., & Turri, J. (2013). Belief through thick and thin. *Noûs*, 47(3).
- Burge, T. (2005). Disjunctivism and perceptual psychology. *Philosophical Topics*, 33(1), 1–78.
- Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
- Burge, T. (2011). Disjunctivism again. *Philosophical Explorations*, 14(1), 43–80.
- Campbell, J. (2002a). Berkeley's puzzle. In T. S. Gendler & J. O. Hawthorne (Eds.), *Conceivability and Possibility*. Oxford: Oxford University Press.
- Campbell, J. (2002b). *Reference and Consciousness*. Oxford University Press.
- Campbell, J. (2011). Relational vs. kantian responses to berkeley's puzzle. In H. L. Naomi Eilan & J. Roessler (Eds.), *Perception, Objectivity and Causation* (pp. 35–50). Oxford: Oxford University Press.
- Campbell, J. (2012). A straightforward solution to berkeley's puzzle. *The Harvard Review of Philosophy*, 18(1), 31–49.
- Cassam, Q. (2011). Tackling berkeley's puzzle. In H. L. N. Eilan & J. Roessler (Eds.), *Understanding Perception, Causation and Objectivity*. Oxford University Press.
- Chisholm, R. (1966). *Theory of Knowledge* (1 ed.). Englewood Cliffs, NJ: Prentice Hall.
- Choi, S. & Fara, M. (2012). Dispositions. *Stanford Encyclopedia of Philosophy*.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. MIT Press.
- Chomsky, N. (1995). Language and nature. *Mind*, 104, 1–61.
- Chomsky, N. (2003). Reply to egan. In L. Antony & N. Hornstein (Eds.), *Chomsky and his Critics*, volume 89-104. Oxford: Blackwell.
- Chrisman, M. (2012). The normative evaluation of belief and the aspectual classification of belief and knowledge attributions. *Journal of Philosophy*, 109(10), 588–612.
- Conee, E. & Feldman, R. (1985). Evidentialism. *Philosophical Studies*, 48, 15–35.
- Dennett, D. (1978). A cure for the common code? In *Brainstorms: Philosophical Essays on Mind and Psychology* (pp. 90–126). Bradford Books.
- Devitt, M. (2006). *Ignorance of Language*. Oxford: Clarendon Press.
- Dickie, I. (2012). Skill before knowledge. *Philosophy and Phenomenological Research*, 85(3), 737–745.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, F. (1986). Misrepresentation. In R. Bogdan (Ed.), *Belief: Form, Content, and Function*. Oxford University Press.
- Egan, F. (1995). Computation and content. *Philosophical Review*, 104, 181–203.

- Egan, F. (2013). How to think about mental content. *Philosophical Studies*, DOI 10.1007/s11098-013-0172-0.
- Evans, G. (1982). *The Varieties of Reference*. New York: Oxford University Press.
- Evans, G. (1985). Understanding demonstratives. In *Collected Papers*. Oxford: Clarendon Press.
- Farah, M. J. (1989). Mechanisms of imagery-perception interaction. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 203–211.
- Field, H. (1978). Mental representation. *Erkenntnis*, 13, 9–61.
- Fodor, J. (1980). Methodological solipsism. In R. Boyd, P. Gasper, & J. D. Trout (Eds.), *The Philosophy of Science (1991)*. (pp. 651–669). MIT Press.
- Fodor, J. (1987). *Psychosemantics*. MIT Press.
- Fodor, J. (1990a). *A Theory of Content and other essays*. MIT/Bradford.
- Fodor, J. (1990b). A theory of content II: The theory. In *A Theory of Content and other essays*. MIT/Bradford.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Ginsborg, H. (2009). Perception, generality and reasons. In A. Reisner & A. Steglich-Petersen (Eds.), *Reasons for Belief*. Synthese Library, Springer Press.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 771–791.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and Knowledge*. Dordrecht: Reidel.
- Greco, J. (2001). Virtues and rules in epistemology. In A. Fairweather & L. Zagzebski (Eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility* (pp. 117–141). Oxford: Oxford University Press.
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies*, 142(1), 17–26.
- Greco, J. (2010). *Achieving Knowledge*. Cambridge University Press.
- Greco, J. (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85(1).
- Greco, J. & Turri, J. (2011). Virtue epistemology. *Stanford Encyclopedia of Philosophy*.
- Hájek, A. (2007). The reference class problem is your problem too. *Synthese*, 156, 185–215.
- Hetherington, S. (1999). Knowing failably. *Journal of Philosophy*, 96(565–587).
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.

- Howe, C. & Purves, D. (2005a). The müller-lyer illusion explained by the statistics of image-source relationships. *Proceedings of the National Academy of Sciences*, 102(4), 1234–1239.
- Howe, C. & Purves, D. (2005b). *Perceiving Geometry*. New York: Springer.
- Jiang, Y. & Chun, M. M. (2001). The spatial gradient of visual masking by object substitution. *Vision Research*, 41, 3121–3131.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24(2), 175–219.
- Kant, I. (1781/1787/1997). *Critique of Pure Reason*. Cambridge and New York: Cambridge University Press.
- Kaplan, D. (1989). Demonstratives.
- Kelp, C. (2013). Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91(2), 265–278.
- Klein, P. (1971). A proposed definition of propositional knowledge. *Journal of Philosophy*, 68(16), 471–482.
- Klein, P. (1976). Knowledge, causality, and defeasibility. *Journal of Philosophy*, 73(29), 792–812.
- Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, 158, 345–361.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies*, 142, 27–42.
- Lehrer, K. (1965). Knowledge, truth, and evidence. *Analysis*, 25, 168–75.
- Lewis, D. (1997). Finkish dispositions. *The Philosophical Quarterly*, 47, 143–158.
- Locke, J. (1975/1690). *An Essay Concerning Human Understanding*. New York: Oxford University Press.
- Luck, S. & Vogel, E. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279–281.
- Lycan, W. (2013). On the gettier problem problem. [http://www.unc.edu/ujanel/Gettier.htm#\\_ednref8](http://www.unc.edu/ujanel/Gettier.htm#_ednref8).
- Martin, M. (2002). Particular thoughts and singular thought. *Royal Institute of Philosophy Supplement*, 51, 173–214.
- Martin, M. (2006). On being alienated. In T. S. Gendler & J. Hawthorne (Eds.), *Perceptual Experience*. Oxford University Press.
- McDowell, J. (1994). *Mind and World*. Harvard University Press.
- McDowell, J. (2010). Tyler burge on disjunctivism. *Philosophical Explorations*, 13(3), 243–255.

- McLaughlin, B. (2010). The representational vs. the relational view of visual experience. *Royal Institute of Philosophy Supplement*, 67.
- Millar, A. (2008). Perceptual-recognitional abilities and perceptual knowledge. In *In Adrian Haddock & Fiona Macpherson (eds.), Disjunctivism: Perception, Action, Knowledge*. Oxford University Press.
- Millikan, R. (1989). Biosemantics. *Journal of Philosophy*, 86, 281–97.
- Miracchi, L. (2014a). Achievements and exercises. In *Getting Things Done (Ph.D. Thesis)* chapter 2. Rutgers University, New Brunswick.
- Miracchi, L. (2014b). Competence to know. In *Getting Things Done (Ph.D. Thesis)* chapter 3. Rutgers University, New Brunswick.
- Miracchi, L. (2014c). A direct virtue aistheology. In *Getting Things Done (Ph.D. Thesis)* chapter 4. Rutgers University, New Brunswick.
- Moore, C., Yantis, S., & Vaughan, B. (1998). Object-based visual selection: Evidence from perceptual completion. *Psychological Science*, 9, 104–110.
- Myers-Schulz, B. & Schwitzgebel, E. (2013). Knowing that p without believing that p. *Noûs*, 47(2), 371–384.
- Noë, A. (2003). Causation and perception: the puzzle unravelled. *Analysis*, 63(2), 93–100.
- Palmer (1999). *Vision Science: From Photons to Phenomenology*. MIT Press.
- Pappas, G. (2007). Berkeley's assessment of locke's epistemology. In S. H. Daniel (Ed.), *Reexamining Berkeley's Philosophy*. University of Toronto Press.
- Pautz, A. (2011). Can disjunctivists explain our access to the sensible world? *Philosophical Issues*, 21, 384–433.
- Plantinga, A. (1993). *Warrant and Proper Function*. Oxford: Oxford University Press.
- Pritchard, D. (2010). Anti-luck virtue epistemology. In A. H. Duncan Pritchard, Alan Millar (Ed.), *The Nature and Value of Knowledge: Three Investigations* chapter 3. Oxford University Press.
- Purves, D. & Lotto, R. (2002). *Why We See What We Do: An Empirical Theory of Vision*. Sunderland, MA: Sinauer Associates.
- Purves, D., Lotto, R., Williams, S., Nundy, S., & Yang, Z. (2001). Why we see things the way we do: Evidence for a wholly empirical strategy of vision. *Philosophical Transactions of the Royal Society of London, B Divison (Biological Sciences)*, 356, 285–297.
- Purves, D., Wojtach, W., & Lotto, R. B. (2011). Understanding vision in wholly empirical terms. *Proceedings of the National Academy of Sciences*, doi:10.1073/pnas.1012178108.

- Putnam, H. (1981). Reprinted as “Brains in a Vat”, in DeRose and Warfield (1999). *Reason, Truth, and History*. Cambridge University Press.
- Radford, C. (1966). Knowledge—by examples. *Analysis*, 27(1–11).
- Riggs, W. (2003). Understanding ‘virtue’ and the virtue of understanding. In M. DePaul & L. Zagzebski (Eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.
- Roberts, R. C. & Wood, W. J. (2007). *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Clarendon Press.
- Robinson, H. (1994). *Perception*. Routledge.
- Rock, I. & Gutman, D. (1981). The effect of inattention on form perception. *Journal of Experimental Psychology: Human Perception and Performance*, 7(275–285).
- Rose, D. & Schaffer, J. (forthcoming). Knowledge entails dispositional belief. *Philosophical Studies*.
- Schellenberg, S. (2010). The particularity and phenomenology of perceptual experience. *Philosophical Studies*, 149, 19–48.
- Schellenberg, S. (2013). Experience and evidence. *Mind*, 122(487), 699–747.
- Scholz, B. C., Pelletier, F. J., & Pullum, G. K. (2011). Philosophy of linguistics. *Stanford Encyclopedia of Philosophy*.
- Searle, J. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences*, 3, 417–57.
- Searle, J. (1984). *Minds, Brains, and Science*, chapter Can Computers Think?, (pp. 28–41). Harvard University Press.
- Sellars, W. (1997/1956). *Empiricism and the Philosophy of Mind*. Harvard University Press.
- Siegel, S. (2004a). Indiscriminability and the phenomenal. *Philosophical Studies*, 120, 90–112.
- Siegel, S. (2004b). Review of *Reference and Consciousness*, by John Campbell. *Philosophical Review*, 113(3), 427–431.
- Siegel, S. (2008). The epistemic conception of hallucination. In F. Macpherson & A. Haddock (Eds.), *Disjunctivism: Perception, Action, Knowledge* (pp. 205–224). Oxford University Press.
- Siegel, S. (2010). *The Contents of Perception*. Oxford University Press.
- Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy*, 5, 3–25.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press.

- Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives* 24: *Epistemology*, 465–475.
- Sosa, E. (forthcoming). The unity of action, perception, and knowledge. *Forthcoming*.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2011). *Know How*. Oxford University Press.
- Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT/Bradford.
- Tipper, S. & Behrmann, M. (1996). Object-centered not scene-based visual neglect. *Journal of Experimental Psychology: Human Perception and Performance*, 22(5), 1261–1278.
- Tipper, S., Driver, J., & Weaver, B. (1991). Object-centered inhibition of return of visual attention. *Quarterly Journal of Experimental Psychology*, 43A(289-298).
- Travis, C. (2013). *Perception: Essays After Frege*, chapter Introduction. Oxford: Oxford University Press.
- Turri, J. (2011). Manifest failure: The gettier problem solved. *Philosophers' Imprint*, 11(8), 1–11.
- van Gelder, T. (1998). The dynamical hypothesis. *Behavioral and Brain Sciences*, 21(5), 615–28.
- Weatherson, B. (2003). What good are counterexamples? *Philosophical Studies*, 115, 1–31.
- Williamson, T. (2000). *Knowledge and Its Limits*. Oxford University Press.
- Wilson, G. & Shpall, S. (2012). Action. *Stanford Encyclopedia of Philosophy*.
- Zagzebski, L. (1994). The inescapability of gettier problems. *Philosophical Quarterly*, 44(174), 65–73.
- Zagzebski, L. (1996). *Virtues of the Mind*. Cambridge: Cambridge University Press.