Making Sense of Hallucination

A THESIS

SUBMITTED TO THE FACULTY OF THE UNIVERSITY OF MINNESOTA BY

LINK SWANSON

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF ARTS IN PHILOSOPHY

Peter Hanks, Advisor

October 2019

Acknowledgements

I recognize Kathryn Swanson, Ph.D, JD, my wonderful philosopher-wife, for inspiring me to enter graduate school and for continually giving me the guidance, advice, encouragement, and support I needed to be successful. Our discussions of the ideas in this work were immensely helpful, and her capacity for deep thinking—both analytical and creative—has rubbed off on me in ways that have truly transformed every aspect of my intellectual pursuits.

I acknowledge Joseph Owens for giving me a very solid 'lay of the land' in philosophy of mind early on in my graduate career, and for inspiring me with his passion for these issues and ideas. I got a solid start, thanks to Joe.

I acknowledge Peter Hanks for jumping off the deep end of both philosophy of perception and Kant with me during a series of memorable graduate seminars. His ability to bring clarity and definition to the process of philosophy is remarkable. Furthermore, his dedication to quality mentorship was a true gift to me as a graduate student, and a large factor in my scholarly success.

I acknowledge Dan Kersten for repeatedly validating my hunch that thinking like a philosopher is a valuable tool in the science of perception. Dan introduced me to predictive coding and its Helmholtzian roots, which was pivotal for me. I am also indebted to Dan for his ongoing mentorship.

I acknowledge the influential conversations I have had with certain brilliant thinkers, including but not limited to Howard Robinson, Galen Strawson, Michelle Montague, Dennis McKenna, Bob Fisk, Jacob Gossel, and Ben Prigge.

I acknowledge my children, Huxley and Juniper Swanson, who ask (and answer) the very best questions about mind and perception. Not only do they remind me of the wonder and mystery in all conscious experience, but they also fill me with love, motivation, and gratitude.

Dedication

For Juniper

Abstract

It can seem as if philosophy of perception has discussed hallucination almost more than perception itself. What is the difference between perception and hallucination? I argue that the concepts we normally associate with the term 'hallucination' are more useful for understanding what perception *is* than the concepts we normally associate with the term 'perception'. Instead of claiming, as most theories do, that hallucination is a special type of (failed) perception, I instead argue that perception is a special type of (successful) hallucination. I introduce a concept called '*S*-hallucination' and argue that it more accurately describes the process that we normally call perception. I defend this concept and situate it within classic debates in philosophy of perception.

Table of Contents

List of figures

The philosopher who would do justice to perception today must have in addition to his philosophical training an advanced knowledge of neurophysiology, electronics, neurology, psychology and psychiatry, and preferably should himself have taken a hallucinogenic drug.

—Sir Russell Brain (1956, p. 2)

Recreational hallucinations were very common in the latter decades of the twentieth century, when many people took "hallucinogenic" drugs. I never did this, but I know people who did.

—John R Searle (2015, p. 163)

... some philosophers are prepared to write at length about hallucinations without ever having made a close study of these phenomena, either by reading the extensive psychopathological literature, or by talking to people who are having or who have had hallucinations, or, best of all, by having some hallucinations for themselves.

—John Smythies (1956, p. 2)

As far as I know, I have never had a hallucination in my life.

—John R Searle (2015, p. 163)

Introduction

Here is the situation in philosophy of perception. In everyday life, our sensory faculties deliver to us scenes of an external world presenting various objects, entities, and events. Right now you likely experience some walls, windows, a table, a chair, and maybe a mug on the table. You can feel the chair beneath you, you can see the page in front of you, you can smell and taste the beverage in your mug. Looking out your window, you maybe see a lavender bush, or some pine trees, or a snow-covered churchyard, or dappled deer grazing in groups. Your experience presents these objects, entities, and events as having an immediate existence in the external world—you can *see them*; they are *there*. They appear to exist independent of your—or anyone else's—seeing them.[1] Describe your experience to me, and you inevitably invoke descriptions of the objects themselves. Introspect, reflect, and philosophize all you want about your experience, but it will still present to you objects, entities and events embedded within a scene that apparently *is* the external world, *right there* before you.

Since our perceptual experiences unavoidably have the phenomenal character of presenting an external world, it is natural to assume that the objects, entities, and events presented in a perceptual scene *really are there* and *really are as they appear*. And by 'natural' I mean that most people, children, and animals intuitively and instinctively trust that what they experience as in the scene *is* the world, and that the objects, entities, and events presented therein *actually have* the properties they appear to have. Indeed, *we must* assume this for obvious immediate safety reasons. Furthermore, most humans, when asked to reflect, will assert that what they experience in the scene *is the world itself* and that they can use their sensory faculties to enjoy direct access to the actual properties of external objects, entities, and events. This intuition is known in philosophy as *direct realism*.

[1] As PF Strawson (1979, p. 97) phrases it, "mature sensible experience (in general) presents itself as, in Kantian phrase, an immediate consciousness of the existence of things outside us."

Another name for direct realism is *naive realism*. Why the pejorative 'naive'? Why would it be naive to believe, for example, that the massive tumbling boulder that you can see, hear, and feel barreling toward you *really is* a large, heavy rock that *really is* tumbling toward you and *really could* crush you? What seems naive—or at least foolish—is to *question* such useful appearances. However, philosophers, especially since the early modern period, have exposed and criticised the *fallibility* of direct realism. Hume, for example, remarked that "this universal and primary opinion of all men is soon destroyed by the slightest philosophy" (Hume, 1777, p. 151). Reid shared Hume's assumptions, stating that "All philosophers, from Plato to Mr. Hume, agree in this, that we do not perceive external objects immediately, and that the immediate object of perception must be some image present to the mind. So far there appears to be a unanimity rarely to be found among philosophers on such abstruse points" (Reid, 1785). While Reid's unanimity claim is a bit much (and historically false), it is true that a general rejection of direct realism can be found in the writings of Descartes, Locke, Berkeley, and Hume—philosophers who, like the ancient[2] and medieval[3] philosophers before them, emphasized the fact that at least *some* of the sensible qualities displayed by the objects in our perceptual experience are *not actually properties possessed by those objects*. For example, a paddle appears bent when half-submerged in water, yet it is not *actually* bent. The size and shape of a table or a coin can change when viewed from different angles or from varying distances, yet this does not *actually* change the size or shape of the coin or the table. Hold up a coin and look at it, then press one finger against your eye and you might get 'double vision' and see two coins. Do you now have twice as much money?

These common optical illusions and aspectual variations of appearances can motivate a legitimate questioning of direct realism. Stronger questioning can be motivated by reflecting on more extreme, yet still very common phenomena, such as dreams, delusions, and hallucinations. Moreover, somatic own-body perception can present illusory phenomena such as phantom limbs, vertigo, or out-of-body experiences.

---

[2] For instance, Heraclitus, Parmenides, Plato, and Aristotle (see Caston, 2015).
[3] For instance, Aquinas and Ockham (see Perler, 2015).

Even technologies like photographs, phonographs, television, and virtual reality have raised doubts about the core assumptions of direct realism.

Howard Robinson (1994, p. 42) summarizes the problem with our direct realist intuitions about perception:

> … our pre-philosophical idea of what it is to perceive the world contains two incompatible elements. One is that we are in direct perceptual contact with the external world, which is taken to imply that one is connected to the world by 'transparent' conscious apprehending which simply puts one in touch with the object as it is [direct realism]. The other is that objects vary in appearance although they are not changing intrinsically, so that we are not in touch with the object as it is intrinsically. Pre-philosophical common sense about perception is, therefore, self-destructive.

This is the core problem that I address in this thesis. You experience your perceptual field as if it discloses objects independent of the activity of your mind. Yet numerous perceptual situations, illusions, hallucinations, and other forms of perceptual experience contradict the assumption that the objects presented in veridical perception are mind-independent. Most philosophical theories of perception try to show either (1) how the activity of your mind has the power to both contact mind-independent objects and generate mind-dependent illusions and hallucinations (sense-data theories and intentionalistic theories), or (2) how veridical perception is a special kind of direct contact with mind-independent objects and illusions and hallucinations are of a fundamentally different and completely unrelated kind of experience (disjunctivism).

Here, I will argue for a different approach. I argue that the concept of 'mind-independent objects'—that is, the idea that there can be sensible objects that have the properties we experience them to have independent of the activity of our minds—is an idea that must be banished from our philosophical theories of perception. Any meaningful sense of the word 'object' in the context of

perceptual experience, I argue, is something with a mandatory a mind-dependent ontology. The notion that you can experience a mind-independent object during any kind of perceptual experience *is a myth*. This technically makes my approach 'anti-realist' and will likely get me indicted for espousing 'idealism' with a charge of discarding any hope of establishing what is 'real' from what is imaginary. However, I reject this charge, and my rejection is backed up by what is a substantial contribution I hope to make with the present treatment: I offer a philosophical theory of what makes a perceptual experience veridical which *does not* invoke criteria involving the requirement that the experience 'contacts' objects with a mind-independent ontology. Equipped with these criteria, I argue that we can finally start making sense of hallucination, illusion, perception, veridicality, and reality.

Chapter 1

The problems of perception

In the Introduction I briefly sketched the problem of perception in philosophy. Direct realist intuitions are ubiquitous across human and animal perceivers, and most perceptual experience itself has a subjective phenomenal character that presents itself as 'real' and 'external' and 'mind-independent'. At odds with this is the fact that all sorts of perceptual experiences can occur which appear to contradict this intuition; namely, illusions and hallucinations. While Chapter 3 contains detailed arguments against direct realism, here I outline the basic principles which result in puzzles and problems in philosophy of perception.

The phenomenal principle

How can we concisely capture the tension that non-veridical perceptual experiences generate for direct realism? Robinson has captured the core of it for us with what he labels the "Phenomenal Principle," which can be stated like so:

> If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality (Robinson 1994, p.32).

So when you look at your canoe paddle half-submerged in the cool waters of the BWCA, it *looks bent*; you see a bent paddle. The *bentness* of the paddle is perceptually presented to you in your subjective perceptual field and therefore, according to the phenomenal principle, there is something of which you are aware which does possess the

sensible quality of being bent. But—and here is the tension between the phenomenal principle and direct realism—the paddle is not actually bent. The property of being bent, of which you are aware, is beared by *something*, but not the paddle itself. Again, direct realism is the assumption that the sensible properties of the objects presented in your subjective perceptual field are the properties the objects actually possess. Yet here the paddle that you clearly see as bent does not actually possess the property of being bent. Thus, there must be another something, of which you are aware, that is bent. If the phenomenal principle is true, then since that other something is not the actual paddle, direct realism must be false.

The problem can be emphasized using a more extreme illustration. Again suppose you look at your canoe paddle half-submerged in the cool waters of the BWCA, but you are actually fast asleep in your bed dreaming about a BWCA canoe trip. Or you look at your canoe paddle half-submerged in the cool waters of the BWCA, but you are actually sitting half-on the seat of your rocking chair fully immersed in a BWCA hallucination. Do you see a paddle? According to the phenomenal principle, if the paddle sensibly appears to you in your dream or your hallucinatory state, then there is something of which you are aware which has the properties of being wood and paddle-shaped. But it is not a 'real' paddle in the way 'real' is used to define direct realism.

Arguments like this (some odd fact about perception as a premise; phenomenal principle as an additional premise; therefore direct realism is false) are the root of the 'problem of perception' defining the past several centuries of philosophy of perception (Crane, 2005; Crane & French, 2017; Hatfield, 2015; Searle, 2015; see Simmons, 2015; Smith, 2002). On the one hand, the phenomenal principle is a powerful intuition. On the other hand, direct realism is also a powerful intuition. Accept the phenomenal principle and direct realism becomes incompatible with the varieties of non-veridical and partially-veridical perceptual experience. By far the most common move has been to accept the phenomenal principle, reject direct realism, then offer a new theory as an alternative to direct realism with a better account of what makes up the varieties of perceptual experience. Though the nature of their accounts differ, this overall approach is

taken up by Descartes, Malebranche, Locke, Hume, Boyle, Berkeley, Reid, and Kant (Hatfield, 2015; see Simmons, 2015).

In the 20th century, philosophers started to carve out labels to differentiate the types of possible responses to the problem of perception. *Sense-data* theorists claim that the objects presented in your perceptual field are made of a special kind of stuff—sense-data. *Intentionalistic* theorists claim that the objects presented in your perceptual field, akin to beliefs or propositions, are *about* the objects in the world. *Disjunctive* theorists, by contrast, attempt to accept both the phenomenal principle and direct realism by questioning assumptions about the facts of hallucination or illusion so as to 'disjoin' these experiences from veridical perception. For example, some disjunctive theorists have claimed that in having your hallucination you are not aware of anything, or that your hallucination simply has no phenomenal character whatsoever (Fish, 2009).

As Robinson (1994) and Searle (2015) point out, philosophers of the modern period from Descartes to Hume to Kant seem to have accepted the phenomenal principle as unavoidable. Moreover, twentieth-century theorists Price, Broad, Moore, and others also held to the principle. Why? As Robinson explains, "the principle is very closely associated with a belief in the self-intimating nature of consciousness. All these philosophers believed that there could be no certainty greater than that given by something's being clear and distinct to consciousness; plain phenomenology is the foundation of everything" (Robinson, 1994, p. 36).

The common-kind thesis

When we accept the phenomenal principle and proceed to analyze the varieties of perceptual experiences, a second question arises. Does your veridical experience of paddling a canoe across the BWCA share any properties in common with your (*ex hypothesi*) identical hallucinatory experience? Since the experiences are identical, we might reasonably claim that the phenomenal character shared between these two

experiences entails a shared fundamental kind of mental state or event. This is known as the *common-kind thesis*[4]: the claim "that the most basic phenomenal character in both experiences is the same" (Siegel, 2008, p. 205). The common-kind thesis is in tension with any theory claiming to demarcate hard ontological boundaries between perception, illusion, and hallucination. This is especially true of direct realism: How can mind-independent objects be essential for the phenomenal character of a veridical perception but not for the hallucination if the common-kind thesis is true? Furthermore, if we accept the common-kind thesis, then we should be able to give a common explanation of what constitutes the phenomenal character they share in common. This leads us into the screening-off problem.

The screening-off problem

What is the best explanation of your hallucination that you are canoeing in the BWCA? Nowadays a common response might be "your brain activity generated the hallucinatory experience" along with the conclusion that the explanatory account should be a neuroscientific one. But what about the explanatory account of your veridical experience of canoeing in the BWCA? A common response to this might be "you experienced canoeing in the BWCA because you really did canoe in the BWCA." The direct realist must include the objects, entities, and events of a BWCA canoe trip as essential components in the explanatory account of your experience. But how can the direct realist claim that these items are 'essential components' of the explanation of your veridical experience when they are not required for the explanation of your identical hallucinatory experience? Any explanation of the hallucinatory case would also explain the veridical

[4] Martin (2004, p. 40) first labeled it as the *Common Kind Assumption*: "whatever kind of mental event occurs when one is veridically perceiving some scene, such as the street scene outside my window, that kind of event can occur whether or not one is perceiving." Fish (2010) labels it the *Common Factor Principle* defined similarly as "Phenomenologically indiscriminable perceptions, hallucinations, and illusions have an underlying mental state in common." Crane and French (2017) phrase it as "Whatever fundamental kind of mental event occurs when one veridically perceives, the very same kind of event could occur were one hallucinating."

experience and so would 'screen off' the direct realist claims as extraneous explanation (Michael G. F. Martin, 2004).

Summary

The situation in philosophy of perception includes the three well-known puzzles described so far. The phenomenal principle is the intuition that when something is presented in your subjective perceptual field, no matter if it is illusory or hallucinatory, there is still *something* which has the phenomenal character presented in your experience. This intuition, combined with facts about common hallucinations and illusions, places pressure on a different intuition; namely, direct realism. Two further puzzles put pressure on direct realist intuitions. First, the common-kind thesis is the claim that the phenomenal character shared between two experiences entails a shared fundamental kind of mental state or event. So your veridical perception of a yellow banana and your non-veridical hallucination of a yellow banana share a common mental state—the experience of seeing a yellow banana. The ontological consequence of this places pressure on the direct realist claim that mind-independent objects are fundamental to whatever kind of mental event underlies veridical perception because you can have the experience of seeing the object whether it is actually there or not. Second, when we analyze the necessary elements of an explanatory account of any given perceptual experience, any explanatory account sufficient to explain the hallucination would also be sufficient to explain the veridical perception and would 'screen off' direct realist assumptions.

Chapter 2

S-Hallucination Theory of Perception

Perception as S-hallucination

Here, I propose a theory of perception which I call the *S-Hallucination Theory of Perception* (SHTOP). My core thesis is captured in the following statement:

> … every "normal" perception may be said to have a "hallucinatory" component … (Klüver, 1966, p. X).

The above quotation is from the preface to the 1966 book *Mescal and Mechanisms of Hallucinations* by Heinrich Klüver[5], a pioneering neurologist of the early twentieth century with a keen interest in hallucination and visual perception. Klüver's statement to captures the gist of my core thesis.

First, the qualifier *every*—"every 'normal' perception"—sets the scope of the claim to cover *all* perceptual experiences, without exception.

Second, the quotations around 'normal'—"every 'normal' perception"—capture a core subclaim which I argue for below: it is not possible to demarcate 'normal' from

[5] Klüver apparently considered this statement to be, in his words, a "fact," as if it would be obvious and uncontroversial to his readers. Given the kind of work Klüver was engaged in (see Nahm & Pribram, 1998), this is understandable.

'non-normal' perception without making arbitrary (and counterproductive) distinctions and assumptions.[6]

Third, the quotations around 'hallucination'—"a 'hallucinatory' component"—capture a further subclaim I argue for below: it is not possible to demarcate 'hallucination' from 'non-hallucination' without making arbitrary (and counterproductive) distinctions and assumptions.

Fourth, the use of the word 'component'—"a 'hallucinatory' component"—captures a *radical common-kind thesis* which I argue for below: Not only are veridical perceptions and hallucinations of the same fundamental kind of mental processes (common-kind thesis); they also coexist together, to different *degrees*, in every perceptual experience (radical common-kind thesis).

Finally, the full claim itself—"every 'normal' perception may be said to have a 'hallucinatory' component"—emphasizes that *all* perception is *always*, in some sense, to some degree, more or less hallucinatory.

My proposal insists on the need for a radical overhaul of the concepts and definitions which underpin the common terms 'perception' and 'illusion' and 'hallucination'. I propose that philosophers begin to think of perceptions in a way that is similar to how we already think about hallucinations. Instead of claiming, as most theories do, that hallucination is a special type of (failed) perception, I instead argue that perception is a special type of (successful) hallucination.

To facilitate this conceptual overhaul, I introduce the technical term *S-hallucination* as a replacement for what philosophers usually call 'veridical perception'—I argue that what goes on in veridical perception is better captured by a modification of our concept of 'hallucination' than it is by our standard concept of perception. For this reason I replace the term 'perception' with the term *S*-hallucination.

[6] Klüver was keenly aware of the great range of diversity in human perceptual faculties as he studied the neurology of visual perception, abnormal vision, hallucination, the effects of mescaline, and he himself took peyote in the Department of Psychology at the University of Minnesota in 1926 (Klüver, 1926, 1928).

Before defining what exactly I mean by *S*-hallucination, I should address what is likely to be an immediate objection to my proposal thus far.[7] Am I simply being dramatic with my language? Am I merely playing tricks with words? Hallucination, for most people, means 'not real' or 'false'—so why am I saying that perception is more like hallucination than it is like our concept of perception? As stated above, I argue that our traditional concept of hallucination offers a more useful way of understanding perception than does our traditional concept of perception. Why should I think our concept of *hallucination* is better suited for what perception *is*? When you think of what a hallucination is, you think of a subjective perceptual experience produced in the mind, or perhaps generated by the brain, independent of anything external to it. For example, let's say that you hallucinate an experience in which you are gently paddling a canoe across the BWCA. You are not actually on a canoe trip; you are hallucinating this while sitting half-on your rocking chair. According to our standard conception of hallucination, the lakes, trees, and mosquitoes presented in your subjective hallucination do not depend on external mind-independent lakes, trees, or mosquitoes. Because we are talking about a hallucination, we easily accept that your mind *somehow* produces the experience endogenously. I argue that this is fundamentally what *all* perceptual experience is. You should think of the contents of your subjective perceptual field as produced in your mind, generated by your brain. But our traditional concept of perception scarcely includes notions of 'produced in the mind' or 'generated by the brain' as part of its connotation. Rather, it predominantly entails quite the opposite, conjuring concepts such as 'seeing things as they are' or 'direct contact' with 'actual properties' of 'external' 'mind-independent' objects. I argue that this connotation of the word *perception* is the root of the philosophical problem of perception. We hold a mistaken concept of what perception *is*, and the word 'perception' itself embodies the concept so strongly that it prevents us from understanding the nature of the actual phenomenon that it is supposed to refer to. The word 'hallucination', on the other hand, is equipped with the concepts needed to understand what perception *is*. This is my rationale for my admittedly

[7] Yoshi Yoshinari, Peter Hanks, and Cat St. Croix urged me to more clearly address this issue.

bombastic proposal of the technical term *S*-hallucination to replace 'veridical perception' in my philosophical theory of what perception is.

What exactly is meant by *S*-hallucination? It it denotes the *high* end of a quasi-quantitative scale of qualities I call *S*-traits, set out as follows:

> *S-hallucination* is any perceptual experience (or *percept*) which exhibits a high degree of the following *S-traits*: (1) *stable*, (2) *successful*, (3) *status-quo*. How are *S*-traits defined? They can be summarized as follows:
>
> *Stable:* The percept occurs in a repeatable manner and remains durable across a variety of perspectives.
>
> *Successful:* The percept enables action or actionable information.
>
> *Status-quo:* The percept is consistent with past experience and with prior knowledge about 'how things are' and 'how the world works'.
>
> *S-factor*: The degree to which a percept exhibits any and all of the above *S*-traits within a given sense modality, or across multiple sense modalities, or across individuals, or across species.

The core claim of *S*-hallucination theory can now be roughly stated. Veridical perception and hallucination can be differentiated only by the degree to which they exhibit the above *S*-traits, which are wholly subjective, phenomenological and thus mind-dependent in their ontology. What philosophers traditionally call veridical perceptions are more appropriately conceived of as hallucinatory phenomena that exhibit a higher *S*-factor; what we traditionally call hallucinations are hallucinatory phenomena that exhibit a low *S*-factor. Importantly, I argue that both phenomena are of the same fundamental ontological

kind. They are both experiences of a subjective perceptual field presenting some phenomena as various combinations of phenomenal character. Perceptual experiences that we tend to call 'veridical perception' exhibit a high degree of *S*-traits—they are more stable, more successful, and more in keeping with the status-quo; they have a high *S*-factor. Perceptual experiences that we tend to call 'hallucination' exhibit a low degree of *S*-traits—they are less stable, less successful, and less status-quo; they have a low *S*-factor. This, I argue, is how we should distinguish veridical perceptions from hallucinations. Perception is *S*-hallucination.

Importantly, I argue that *S*-factor, or something like it, is the *only* legitimate way to differentiate 'veridical perceptions' from hallucinations or illusions. Most theories—from direct realism to disjunctivism to intentionalism to sense-data relationalism to representative realism—attempt to disjoin veridical perception from hallucination using various acrobatic moves hoping to demonstrate how their account can give veridical perception an exclusive relationship status with 'mind-independent' objects. I argue that this is a mistake and that notions of mind-independent objects have no place in philosophical theories of perception. This point will be visited repeatedly in the arguments contained in subsequent chapters.

But, you might say, it sounds very harsh to say that the objects in perception are *S*-hallucinations, that we eat and drink *S*-hallucinations, and are clothed with *S*-hallucinations. I acknowledge this—the word *hallucination* not being used in common discourse to signify the several combinations of phenomenal qualities which are commonly called *perceptions*. These expressions depart from the familiar use of language and thus will certainly seem harsh and ridiculous. But this does not change the fact that the usual connotations of the word 'perception' fail to deliver the concepts needed to understand what perceptual experience *is*. I agree that it would be more proper or conformable to custom to call it perception rather than *S*-hallucination. Part of my project is to induce a productive conceptual overhaul so that the word *perception* can eventually

be used in such a way that its connotations include the concepts necessary for understanding what perception actually is.[8]

Another objection is inevitable at this point. Am I banishing all that is real and substantial in nature to some chimerical scheme of hallucinations?[9] What are we to think of houses, rivers, mountains, trees, and stones—or even our own bodies? Are these objects mere hallucinatory chimeras, illusions on the fancy? Am I giving up on veridicality? No. Quite the opposite—I am salvaging it from the puzzles posed by the mistaken assumptions which created the problems of perception in philosophy. In spite of the fact that I argue that perception is *S*-hallucination, I nonetheless maintain that the 'reality' of your veridical experiences—what you see or feel or hear—remains as secure as ever, and is as real as ever. On the principles of the *S*-traits, the distinction between reality and fantasy retains its full force. Indeed, as I will argue in the chapters to follow, the *S*-factor scale provides a robust method of distinguishing how much 'reality' is 'in' a given perceptual experience. Furthermore, since the *S*-traits are rooted *solely* in attributes of the phenomenal character of the experience (as it unfolds across time within the context of other experiences) this affords a built-in feature: definitions of how much 'reality' is 'in' a given experience are not only possible using *S*-factor scale, but are uniquely resistant to sceptical attacks.

In summary, I am arguing that perception and hallucination are at bottom the same fundamental kind of event or process and that the connotations we commonly attach to the word *hallucination* uniquely provide the necessary concepts for understanding what perception *is*. The way to distinguish them is *not* to hunt for links with 'mind-independent' objects, because I argue that the objects of all veridical perceptions are *just as mind-dependent* as the objects of a total hallucination. Rather, I have introduced a systematic metric—*S*-traits and *S*-factor—which can meaningfully assess 'how much reality' is 'in' a given experience and therefore determine which experiences deserve the label of veridicality. I argue that we can accomplish all of this without invoking mind-independent ontology.

---

[8] This paragraph is a tribute to Berkeley (1710, pt. 1, sec. 38) and is mostly in his own words.
[9] This paragraph is also a tribute Berkeley (1710, pts. 1, sec. 34–35).

A further and interesting aspect of my theory must be made explicit. *S*-factor is also a plausible theory of how the mind itself naturally 'makes sense' of its own phenomenal perceptual patterns. The *S*-traits—stability, success, and status-quo—can bootstrap each other to allow the mind to produce an experience that 'feels real' and by my theory actually is real. The general idea can be stated as follows. First, certain percepts can lead to certain other desired percepts; this is my definition of the *S*-trait 'success'. A successful percept eventually becomes habitual and this makes it a stable percept, and stable percepts aggregate to become a status-quo set of expectations about the world. The idea is that the mind bootstraps its own phenomenal framework for classifying veridical vs non-veridical phenomenal patterns. Something like *S*-factor plausibly guides the production of perceptual experience that *feel* veridical.[10] It makes sense, then, that *S*-factor would also appeal as a philosophical basis for veridicality. I argue that we intellectually prosper when we use this metric to understand both how the mind produces perceptual experiences that *feel* real, and how to establish veridicality in our philosophical theory of perceptual experience.

Historical precedent

The general approach of *S*-hallucination theory is not without historical precedent. Klüver apparently ascribed to its core claims as part of his working understanding of visual experience. In the early 1970's a provocative phrase emerged in computer science circles: "Vision is controlled hallucination." This catchy slogan, widely attributed to Max Clowes [11] (Bornat & Brady, 1976; Horn, 1978), became a sort of esoteric adage used by select computer scientists and psychologists working on computer vision (Dayan, Hinton, Neal, & Zemel, 1995; Grush, 2004; Horn, 1980; Neumann, 2003; Sloman, 1984; Zucker, 1987). In the context of this literature, the term 'hallucination' refers to the computer's

[10] I provide a more detailed technical account in the Appendix.
[11] Bornat and Brady (1976) reference the following citation as the source of this phrase:
Clowes, M. B. (1973). Draft notes on vision (privately circulated). University of Sussex.

own internal simulation/emulation of its environment, which it must *constrain* or 'control' in order to successfully 'perceive' the world. With this strategy, the computer uses its stored prior knowledge, expectation, and internal models of the world to 'dream' or 'imagine' or 'hallucinate' an internally-generated environment as a means to perceive its actual physical environment.[12]

Thomas Metzinger has articulated the philosophical implications of the idea in his (2003) book, arguing that philosophy could benefit from the simulation/hallucination description of perception:

> [A] fruitful way of looking at the human brain, therefore, is as a system which, *even in ordinary waking states, constantly hallucinates at the world*, as a system that constantly lets its internal autonomous simulational dynamics collide with the ongoing flow of sensory input, vigorously *dreaming at the world* and thereby *generating the content of phenomenal experience* (Metzinger, 2003, p. 52; emphasis mine).

Metzinger's account is particularly relevant to philosophy of mind and perception because of its claim that *all* perception leverages internal simulations, akin to hallucination or dreaming, to *generate* phenomenal experience—even during ordinary waking veridical perception of the world. Cognitive

---

[12] "We view visual perception as a process of constructing an interpretation of a picture of a scene, not only from visual 'cues' within the picture but also by using knowledge about the ways in which such scenes can be put together. Thus, following Helmholtz, we believe that the order which we perceive is not completely contained in the visual data but is imposed upon it by our preconceptions of what the scene represents. Clowes (1973) has called this 'controlled hallucination'" (Bornat & Brady, 1976, p. 14).

"More recently, discouraged by the complexity of the distributions of raw image intensities, researchers have turned to methods which exploit prior knowledge about the likely contents of the scene being viewed (Reddy et al 1973, Tenenbaum & Barrow 1976). In Max Clowes' words: 'Vision is controlled hallucination'. The image contributes a small 'controlling' influence on the vision system's "hallucinations" based on expectations and predictions" (Horn, 1978, p. 7).

"Another shift in emphasis suggested by this account is that perception is shown to be not a matter of starting with materials provided in sensation and filling in blanks until a completed percept is available. Rather, completed percepts of the environment are the starting point, in that the emulator always has a potentially self-contained environment emulator estimate up and running. This self-contained estimate is operational not only during imagery, but presumably also during dreaming (see Llinas & Pare 1991). The role played by sensation is to constrain the configuration and evolution of this representation. In motto form, perception is a controlled hallucination process" (Grush, 2004, p. 393).

neuroscientist Chris Frith reached a similar conclusion, "You could say that our perceptions are fantasies that coincide with reality" (Frith, 2007, p. 135). Andy Clark cautiously endorses the 'controlled hallucination' metaphor but argues that it might instead be better "to regard hallucination as a form of 'uncontrolled perception'" (2015, p. 308 n.3). Recently, cognitive neuroscientist Anil Seth has shamelessly promoted the phrase in popular media.[13] "I love this phrase," says Seth in an interview with NPR, "I wish I could take credit for it, but I can't." (Raz & Seth, 2018).

The most obvious historical precedent for this line of thinking is the work of Berkeley (1709, 1710, 1713). Berkeley argued that all objects, entities, and events presented in your perceptual field are mind-*dependent*. All you ever directly experience, whether real or imaginary, are what he called 'ideas' but are better defined simply as shifting patterns of subjective phenomenal properties. One of Berkeley's central claims, for which he gives several arguments, is the claim that any notion of 'mind-independent objects' is untenable: in order for something to be an object, Berkeley argues, it must be *perceived as* an object. Therefore all objects are mind-dependent. To serve as a constant reminder of this, Berkeley proposes a shift in terminology to facilitate an overhaul of conceptual habits.

A common objection to Berkeley's idealist ontology is that it *prima facie* appears to collapse any demarcation between the *real* and the *imaginary*. Berkeley vehemently rejects this characterization and reassures us that "the distinction between realities and chimeras retains its full force … That the things I see with mine eyes and touch with my hands do exist, really exist, I make no question" (Berkeley, 1710, pts. 1, sec. 34–35). How does Berkeley secure reality in full force? He invokes criteria very similar to *S*-traits. "The ideas of sense are agreed to have more reality in them—i.e. to be more *strong, orderly, and*

[13] In spite of this disclaimer, Seth has recently touted the slogan as if it were his own in a slew of appearances in popular media, including *TED* (Seth, 2017), *National Public Radio* (Raz & Seth, 2018), *Motherboard* (Rao, 2018), *The Guardian* (Glaser & Sanderson, 2018), and even *Business Insider* (Atalla, Orwig, & Salter, 2018) of all places.

coherent—than ideas made by the mind; but this doesn't show that they exist outside the mind" (Berkeley, 1710, pt. 1, sec. 33).

Apart from Berkeley, *S*-hallucination theory recalls more general philosophical theories about the role of *imagination* in perception (see Kind, 2016). Medieval philosophers discussed the notion of *phantasia* (imagination) as a requirement for perception.[14] Hume's 'faculty of the imagination' was central to his theory of perception (Hume, 1739). Kant claimed quite explicitly that "the imagination is a necessary ingredient of perception itself" (Kant, 1787/1996, sec. A 120 n) and this discussion was revived by Strawson (1970) in the twentieth century. Interestingly, comparisons between the theories of Hume and Kant and the recent theories on which 'vision is controlled hallucination' can be found in Flores (2015) and Swanson (2016) respectively. For now we can simply note that philosophers have long considered that perception requires, in Kant's words, the "productive function of the imagination" (Kant, 1787/1996, sec. A 120–122) and that the two differ, in Hume's words, "only in their strength and vivacity" (Hume, 1739, sec. 1.1.7.5).

However, *illusions*, unlike imagination, have the *same* "degree of force and vivacity" as perceptions, and can be just as orderly. Subject-matter experts on optical illusions—from Helmholtz (1867) to Richard Gregory (1968) to Dale Purves (2003; 2015) to Donald Hoffman (2012; 2014)—continuously insist that illusions and perception come as a packaged deal. Hoffman has called for a revised definition of hallucination to better reflect the facts about the ubiquitous illusory features of everyday perception. "A hallucination *is a perception* experienced by few people that occurs in the absence of an appropriate context and that is not an adaptive guide to behavior" (Hoffman, 2012, p. 14; emphasis mine.).

[14] "all the properties apprehended by the five external senses come together in this internal sense. Their connection and unification enables the faculty of imagination (*phantasia*), another internal sense, to come up with a 'phantasm', i.e. a sensory image that presents a unified whole" (Perler, 2015, p. 3).

The lexicon and conceptual framework I am proposing here answers Hoffman's call for a conceptual overhaul of what is meant by perception, illusion, and hallucination. Hoffman's revised definition, rendered in my terminology, states that hallucinations *are perceptions* which exhibit relatively less stability, successfulness, and status-quo; they have a low *S*-factor.

Computational neuroscientist Cyriel M. A. Pennartz also shares the spirit of *S*-hallucination theory in his analysis of representational systems. "In conclusion, for a well-behaved representational system, it does not matter whether a neural activation pattern represents an external world feature *truthfully* or not, as long as the system represents the feature *consistently* over time and in sensory space" (Pennartz, 2015, p. 279).

The *S*-factor scale is intended to be a metric by which our subjective experiences can be evaluated for veridicality using criteria that do not depend on linking the experience to objects that have a mind-independent ontology. Berkeley proposed a similar solution when he claimed that veridical perceptions are more "strong, orderly, and coherent" than non-veridical perceptions. C.D. Broad argued that we already use something akin to my 'status-quo' *S*-trait, and my *S*-factor scale in general, stating that we evaluate the veridicality of an experience "by considering the inter-relations of experiences with the earlier and later experiences of the same person and the contemporary experiences of others" (Broad, 1952, p. 12).

A theory for 21st-century philosophers

The previous section chronicled a handful of capable thinkers who have in various ways proposed a revision—and a blurring—of the traditional conceptual boundaries surrounding perception, illusion, and hallucination. Note that the majority of these proposals came from neuroscientists, computer scientists, and psychologists who have, understandably, not worked out the philosophical implications of their syncretic revisionary positions on the varieties of perceptual experience. Of the philosophers who

have recently explored the 'perception is controlled hallucination' line of thinking—e.g., Hohwy (2013), Clark (2015)—most have offered expository treatments that do not address how the idea would speak to the traditional questions in the philosophy of perception, such as the debates regarding naive realism, disjunctivism, or intentionality and the contents of perception.[15] Furthermore, a systematic metric for distinguishing 'controlled hallucination' from 'uncontrolled hallucination' has not been attempted.

My original contribution with *S*-hallucination theory aims to provide a methodical philosophical account of the implications of this kind of radical revisionary conceptual overhaul of the boundaries of perception, illusion, and hallucination. Furthermore, I hope to go beyond mere 'implications' and offer a workable technical framework that can be utilized by others to make real progress on the perennial problems posed by the puzzles of perceptual phenomenology.

Screening-off and the common kind thesis

Equipped with the rough outline of the *S*-hallucination theory of perception, we can now visit the common-kind thesis. The common-kind thesis, as described by Crane and French (2017), claims that "Whatever fundamental kind of mental event occurs when one veridically perceives, the very same kind of event could occur were one hallucinating." It is true that with *S*-hallucination theory I do, in some sense, embrace the common-kind thesis. This much should be obvious. And, as mentioned above, *S*-hallucination theory goes a step further to propose what appears to be a *radical* common-kind thesis, fully embracing the claim that perceptions and hallucinations are the same *fundamental* kind of mental event. Importantly, however, I also claim that veridical perception (*S*-hallucination) and hallucination, while at bottom the same fundamental kind, can in fact be *classified according to their S-traits* and meaningfully differentiated using the provided *S*-factor metric. *S*-hallucination (aka veridical perception) can be said to be of a

[15] Metzinger (2003) being the exception.

different class because it exhibits a higher *S*-factor than hallucination. This allows us to maintain a meaningful sense of what is 'real' against what is 'imaginary' without invoking talk of 'external mind-independent objects' or whatever. We can embrace the common-kind thesis—even champion a radical version of it—without worrying about the 'spreading step' (P. F. Snowdon, 1980).

This does not mean that the 'screening off problem' is altogether avoided. When it comes to *explanation* of the processes responsible for an *S*-hallucination experience, it still seems that our explanation of the hallucination would also 'screen-off' the *S*-hallucination. For now I can reply that to address this challenge I would need to provide an explanation for *how* an experience comes to have *S*-traits at all. I do this in Chapter 3.

S-hallucination and the phenomenal principle

How does the phenomenal principle play into *S*-hallucination theory? Recall the principle:

> If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality (Robinson 1994, p.32).

I argue that *S*-hallucination theory diffuses the challenges posed by phenomenal principle, for two reasons. First, because I insist that the term 'hallucination' provides a more appropriate conceptual apparatus for understanding what perceptual experience *is* (and I claim that the concepts implied by the term 'perception' only confuse us) the basic principle of the phenomenal principle is transformed. In hallucination, there can appear to a subject to be something which possesses a particular sensible quality without there being some*thing* of which the subject is aware which possesses that sensible quality.

Second, and relatedly, the phenomenal principle assumes what has been labeled an "act-object" structure of perception, which "posits a distinction between the mental act of sensing, and the object which is sensed" (Crane & French, 2017). What happens to the act-object distinction when we render it through *S*-hallucination theory? I have stressed that our general concept of what hallucination is—specifically the concept that the objects in your hallucinatory perceptual field are *produced in your mind* or *generated by your brain*—is a better match for what perception *is* than is the common connotation of the term 'perception'. With the general concept of hallucination, then, it is natural to think that the *mental act* of hallucination *produces the objects presented in* the subjective perceptual field. Thus, when veridical perception is taken to be *S*-hallucination, the act-object distinction is transformed into the distinction between the process that produces an object and the object which is produced. The mental act of sensing *generates* the object which is sensed.

The consequence of this is that the phenomenal principle becomes less of a puzzle. More can and will be said (in Chapters 3 and 4) about the effects that the *S*-hallucination conception of perception has on the phenomenal principle, act-object distinction, and related lines of thinking.

Making sense of real hallucinations

Apart from the lingering feeling of an almost paradox-like awkwardness of labeling veridical perceptions '*S*-hallucinations', a more somber worry looms. The worry is that *S*-hallucination is hopelessly anti-realist. If the objects of perception are always produced in the mind, the worry goes, then the world presented in our perceptions is not a real world. Reality is either reduced to mere "illusions on the fancy" (Berkeley, 1710, pt. 1, sec. 34) or the objects of perception are only appearances and not things as they are "in themselves" (Kant, 1787/1996).

I can provide a very general argument to soothe this worry. My argument is that mind-independent objects do not secure realism and in fact have the opposite effect because hallucination—a very real phenomenon that actually happens to actual people—is then rendered completely mysterious. This mystery opens the door to arguments against direct realism and scepticism about the reality of the world. Stated formally:

(1) If veridical perception requires mind-independent objects, then hallucination is mysterious.
(2) If hallucination is mysterious, then realism about perceptual experience is vulnerable.
Therefore,
(3) The claim that veridical perception requires mind-independent objects undermines realism about perceptual experience.

By contrast, if we can achieve a theory that can make sense of hallucination and not render it mysterious, a theory comfortable with the idea that there can be 'degrees of veridicality' in any given perceptual experience, a theory that can say how some perceptual experiences can have 'more reality' 'in them' than others, a theory that can give a meaningful and substantial account of the varieties of perceptual experience—then this is the theory that stands the best chance to secure a robust realism about perceptual experience that is not vulnerable to well-known attacks. *S*-hallucination theory is my contribution toward such a theory.

Summary

I am proposing a philosophical framework—*S*-hallucination theory—that embraces a radical revision of our traditional concepts of perception, hallucination, and illusion. So radical, in fact, that I label veridical perceptions '*S*-hallucinations' because I argue that

the concepts connected to the term 'hallucination' better equip us to understand what perception *is*. I take Klüver seriously when he states that "every 'normal' perception may be said to have a 'hallucinatory' component" and I echo Berkeley in holding that the objects of perceptual experience "exist nowhere without the mind" (Berkeley, 1710, pt. 1, sec. 11). Importantly, I do not make the claim that all perceptual experience is mere fantasy, reduced to "illusions on the fancy." In the spirit of Berkeley, I offer a metric—the *S*-factor—for making sense of degrees of veridicality in any perceptual experience, including illusions and hallucinations. *S* is for the *S*-traits, defined in this chapter, that a given perceptual experience can exhibit: it can be more or less stable, more or less successful, more or less status-quo. These traits are rooted *solely* in attributes of the subjective phenomenal character of the experience (as it unfolds across time within the context of other experiences) and do not depend on 'mind-independent objects' to establish their veridicality. The degree to which the experience exhibits these *S*-traits determines its *S*-factor. 'Real' perceptions (*S*-hallucinations) have a high *S*-factor. Hallucinations have a low *S*-factor. Nonetheless, *all* of the objects presented in your subjective perceptual field are produced in your mind, generated by your brain. From this premise our task is to make sense of how some experiences can 'have more reality in them' than others.

Elements of Berkeley, Hume, and Kant make up the historical antecedents of my proposal. I also share the sentiments which began brewing in the twentieth century from remarks like that of Klüver, the adage of computer vision hackers that 'vision is controlled hallucination', the 'perception is unconscious inference' tale of the Helmholtz-Bayes predictive coding movement, and the wisdom of those who research optical illusion. Contemporary philosophy of perception, I argue, needs to integrate this stuff. The goal of the present proposal is to jumpstart that process.

Chapter 3

Classical arguments against direct realism

In the previous chapter I claimed that all of the objects presented in your perceptual field, veridical or not, are always produced in the mind, generated by the brain, and, furthermore, that there is no coherent way to consider the objects presented in your perceptual field as mind-independent objects. I use this claim as rationale for my admittedly dramatic proposal to use the term *S-hallucination* to refer to what is usually called 'veridical perception'—I argue that our common use of the word *perception* carries too many connotations bound up with concepts like 'mind-independent objects', and, that the word *hallucination*, by contrast, is already equipped with many of the concepts which I claim are necessary for understanding what perception *is*, because we have no difficulty thinking of hallucination as a subjective experience which presents perceptual objects but which is endogenously produced in the mind, generated by the brain.

I still need to provide arguments to justify advocating a drastic conceptual overhaul of the idea of perception. I need to (1) provide arguments that the objects presented in your subjective perceptual field cannot be coherently claimed to be mind-independent objects; and (2) show how my claim that perception is more like hallucination can do useful philosophical work toward resolving some of the puzzles of perceptual experience. In this chapter, I achieve (1) simply by narrating a brief exposition of the classic philosophical arguments against direct realism. Along the way, I begin to make the case for (2), a case which I continue to develop by comparing my theory to other theories in Chapter 4.

The argument from perspectival variation

Hume (1777) points out that if you look at an object—a table, for example—it seems to you, on an intuitive level, that your perceptual experience places you in direct contact with the actual size and shape of the table (direct realism). But what happens when you move some distance away and view the same table from a slightly different angle? "The table, which we see, seems to diminish, as we remove farther from it: but the real table, which exists independent of us, suffers no alteration …" (Hume, 1777, sec. XII.1). Russell (1912) expands on this example, pointing out that your perceptual experiences of the table's shape, size, color, texture, hardness, and acoustics will vary depending on your perspectival situation. "Thus it becomes evident that the real table, if there is one, is not the same as what we immediately experience by sight or touch or hearing" (Russell, 1912, p. 11).

Phenomena of perspectival variation were used extensively by philosophers prior to the early modern period (ancient Greeks and Sceptics) to argue epistemological points broadly to the effect that knowledge from the senses cannot be trusted. Epistemic considerations in this vein also feature in the discussions of perspectival variation found in Descartes, Locke, Malebranche, Berkeley, Hume, and Reid. Importantly, however, beginning in the early modern period, the phenomena of perspectival variation raised a different question, a metaphysical question of the *ontology* of the sensible qualities presented in perceptual experience: If your perceptual experience of the table presents sensible qualities which diverge from the 'real' table depending on your perspective, then the table in your experience cannot be identical with the table as it is independent of your mind, as the presentations have differing properties. So *what is* the thing which has the perspectival-dependent table qualities presented in your experience? According to Hume " … it [is], therefore, nothing but its image, which was present to the mind. … the existences, which we consider, when we say, *this house* and *that tree*, are nothing but perceptions in the mind, and fleeting copies or representations of other existences, which remain uniform and independent" (Hume, 1777, sec. XII.1). Here Hume uses the pervasive phenomenon of perspectival variation to make a metaphysical claim about the ontology of the properties presented in perceptual experience; namely, he claims that they

cannot be external mind-independent objects and so must instead be internal mental entities.

The argument from perspectival variation put formally is:

1. In perspectival variation, an object presented in your perceptual field can change in size or shape depending on viewing angle.
2. A truly mind-independent object would not actually change size or shape depending on the perspective of an observer.
3. Therefore, the thing that changes shape in your perceptual field is not a mind-independent object.

The conclusion (3) is reached by Leibniz's law—for any x and y, if x is identical to y, then x and y have all the same properties. Since direct realism rejects (3), it must also reject either (1) or (2). But this will be awkward. Rejecting (2) requires an assertion that a mind-independent object itself changes in tandem with your perspectival shift, which undermines the meaning of 'mind-independent'. Does the mind-independent table itself actually change size or shape as you move about the room? If you answer no, then you accept (2). Direct realism can accept (2) only if it rejects (1). But (1) states a basic fact about everyday perceptual experience and can be easily empirically demonstrated. To deny (1) is to deny that mind-independent objects can appear different depending on viewing angle, a claim that seems plainly false given everyday perceptual phenomenology.

For Hume, Locke, Berkeley, and other early modern empiricists, arguments from perspectival variation were compelling enough to convince them that direct realism was false. Naturally they asked: What *are* the objects in your perceptual field? These thinkers typically concluded that the 'objects' of perceptual experience are 'impressions' or 'ideas' or 'images'—objects produced in the mind which therefore cannot be said to be mind-independent. Sense data theory and other theories aimed at metaphysically characterizing the ontology of these mind-dependent objects of perception are reviewed in Chapter 4.

As we will see throughout this section, perspectival variation shares aspects with other perceptual phenomena which can also be used as premise (1) in arguments with a

similar form—all of which put pressure on direct realism. Traditionally the argument from perspectival variation is used to motivate a rejection of direct realism, from which point philosophers typically offer an alternative way of conceiving how perceptual experiences are *related* to external mind-independent objects.

Since I reject any invocation of mind-independent objects, *S*-hallucination theory handles the argument from perspectival variation in the following way. When you look at the table, your perception of the table is produced in your mind, generated by your brain. Thus, when the phenomenal properties of the table vary as you shift perspective (it looks round from above but oval from the side; brown under fluorescent lighting but greenish in the sunlight; large from up close but small from far away), *all* of these properties belong to the object, because the object is produced in your mind. What about the 'real' properties of the table? Is it 'really' round or oval? Rendered according to the *S*-factor metric, if it the table's shape appears to you as round under most conditions (stable), and seeing it as round enables you to successfully manipulate and interact with it (successful), and being round is consistent with your past experience and with that of other perceiving beings (status-quo)—then the table 'really is' round. Seeing it as round scores a high *S*-factor, so the table has the property of being round. Counterfactually, if one person one day walks up to the table and sees it for a moment as triangular, but most days he does not, and everyone else sees it as round, then his experience has a low *S*-factor and we can say that the table is not 'really' round.

The argument from secondary qualities

Our perceptual experiences present various features which include extension, size, shape, motion, sound, color, smell, taste, warmth and cold. Furthermore, these qualities are presented in our experience as a single *Gestalt* field. Historically, some philosophers have argued that it is necessary to carve the perceptible qualities of this field into two sharply distinct metaphysical categories, known as *primary qualities* and *secondary qualities*.

Primary qualities include extension, size, shape, motion and rest. Secondary qualities include color, smell, taste, sound, and tactile feelings such as warmth or cold. What justifies the distinction between primary and secondary qualities? Interestingly, the distinction has historically been motivated by broader metaphysical positions, beginning with the atomist-materialist philosopher-scientists of ancient Greece, such as Democritus and Lucretius. Why did atomist-materialists argue for the primary-secondary quality distinction? Because if matter is supposedly mind-independent, then material objects cannot be said to possess secondary qualities, because these qualities appear to be unavoidably mind-dependent. Making the distinction allowed the materialist-atomists to argue that the primary qualities we know from perception are intrinsic properties of mind-independent matter.

With the emergence of mechanical science in the seventeenth century, Galileo, Descartes, Hobbes, and others critiqued the relationship between 'mechanical' qualities and 'non-mechanical' qualities. One of the central axioms of the mechanical philosophy was that a small set of 'fundamental' mechanical qualities—size, shape, motion and rest—are intrinsic to matter and can produce and explain all other non-mechanical qualities of the universe. Furthermore, the mechanical qualities are perceptible through both vision and touch. The rest of the perceptible qualities—color, sound, smell, taste, warmth and cold—were considered to be among the non-mechanical qualities of the universe. Locke found it useful to label perceptible mechanical qualities as 'primary' and perceptible non-mechanical qualities as 'secondary'—though Locke did not invent the distinction as he was merely labeling a distinction that was already emerging as common practice among his contemporaries. The crucial assumption here is that primary qualities are intrinsic properties of physical objects, independent of perceivers. Secondary qualities, by contrast, are not intrinsic to material objects, and, "so far as their objective existence is concerned, are nothing but mere names for something which resides exclusively in our sensitive body, so that if the perceiving creatures were removed, all of these qualities would be annihilated and abolished from existence" (Galilei, 1623). Galileo's statement, which is a metaphysical claim about the ontological status of

secondary qualities, was shared by Boyle, Descartes, Locke, and others. Primary qualities are the only sensible qualities which directly present intrinsic properties of physical objects; in the words of Locke, "their patterns really do exist in the bodies themselves" (Locke, 1689, secs. II, viii). By contrast, secondary qualities are perceiver-dependent and are "merely certain sensations which exist in my thought" (Descartes, Cottingham, Stoothoff, & Murdoch, 2008 replies to the sixth set of objections). Hume is explicit: "Sounds, colors, heat and cold, according to modern philosophy are not qualities in objects, but perceptions in the mind" (Hume, 1739, bk. III, pt.1 sec. 1).

The first argument from secondary qualities against mind-independent objects of perception can be thus stated:

1. Objects in your subjective perceptual field present secondary qualities.
2. Secondary qualities cannot be said to be mind-independent.

   Therefore,
3. The objects in presented your subjective perceptual field are mind-dependent insofar as they present secondary qualities.

As stated here, the conclusion (3) contradicts direct realism, which "holds that in normal perception we are aware *only* of mind-independent physical objects and their mind-independent properties" (Robinson, 1996; emphasis mine). A familiar pattern emerges. To reject (3) we must deny (1) or (2). Denying (1) will be awkward as secondary qualities are basic and pervasive feature of perceptual phenomenology—colors, sounds, smells, tastes, and various tactile sensations are an everyday part of perceptual experience. Denying (2) is exceedingly problematic, as it is vulnerable to attacks rooted in the basic everyday phenomenological facts of perspectival variation.

However, while (3) contradicts direct realism, it leaves open the possibility of 'primary quality direct realism' on which only the *primary qualities* presented by the objects in your subjective perceptual field are actual mind-independent properties of external objects. But why should qualities of size, shape, motion and rest be considered

mind-independent and not the color, smell, taste, sound, and tactile feelings? Justifying the distinction in a non-arbitrary way without begging the question is challenging.

The "representative realist" gives up direct realism but attempts to maintain the mind-independence of primary qualities by claiming that the primary qualities of the objects presented in your perceptual field *resemble* or *represent* the actual mind-independent properties of the external object. For any red, round object the *roundness* is mind-independent but the *redness* is produced in the mind. Furthermore, this view is often accompanied by the claim that the mind-dependent secondary qualities are *caused by* the mind-independent properties with which primary qualities have an exclusive relationship status.

Berkeley attacks primary-secondary distinction, primary quality direct realism, representative realism, and the doctrine that secondary qualities are caused by mind-independent properties of external objects. First, Berkeley argues that our concepts of size, shape, motion, and rest *depend on* our subjective perceptual experience of their primary qualities, so the claim that they resemble mind-independent properties yields inconceivable entities (Berkeley, 1710, pt. sec. 10–11). Next, he argues that all of the arguments used to show that secondary qualities are mind-dependent also apply to primary qualities or "all other sensible qualities whatsoever" (Berkeley, 1710, pt. sec. 14–15).

With S-hallucination theory I align with Berkeley's arguments against the primary-secondary quality distinction. It is baked into *S*-hallucination theory that so-called primary qualities are mind-dependent in exactly the same way that secondary qualities are widely acknowledged to be. For an object to have extension, be in motion or at rest, or be any particular size or shape, it requires a perceiving mind, because mind is what enables an object to be an object *tout court*. An object bears certain qualities—whether primary or secondary—only when an experience of that object having those properties occurs with a high *S*-factor.

The argument from causal chains

It is a generally accepted fact that our perceptual experiences result from some kind of causal processes. The causal chain begins with physical disturbances in the external environment, which in turn produce physical disturbances at your sense organs, which in turn leads to a chain of physiological events terminating somewhere in your brain, resulting in a perceptual experience. Descartes, Locke, and other empiricists assumed that sensory processes require such a causal chain; indeed, it was implicit in their views on secondary qualities as they claimed that certain arrangements of primary qualities have the 'disposition' or 'power' to produce secondary qualities in perceiving creatures. Berkeley used this assumption in an explicit argument against direct realism:

> No-one can deny the following to be possible: A thinking being might, without the help of external bodies, be affected with the same series of sensations or ideas that you have, imprinted in the same order and with similar vividness in his mind. If that happened, wouldn't that thinking being have all the reason to believe 'There are corporeal substances that are represented by my ideas and cause them in my mind' that you can possibly have for believing the same thing? Of course he would; and that consideration is enough, all on its own, to make any reasonable person suspect the strength of whatever arguments he may think he has for the existence of bodies outside the mind (Berkeley, 1710, sec. 20).

Causal chains of sensory activity have been leveraged in an argument against direct realism:

1. Perceptual experience presumably results from a chain of causal links spanning from the external object to your sense organs and finally terminating in your brain.
2. The object in your perceptual field is directly caused by the final link in the causal chain, namely, some neurobiological event in your brain.
3. Opposite ends of a causal chain cannot be identical.

   Therefore:

4. The object in your perceptual field is not the external object itself.

The conclusion (4) contradicts direct realism—it states that the object in your perceptual experience must be something other than the external object itself. Importantly, this argument must establish that the *object* in your perceptual field is not the external object itself, otherwise the conclusion would merely state that your *experience* is not the object itself, and nobody wants to claim that an experience and the object of experience are identical (Robinson, 1994, pp. 84–87; although see G. Strawson, 2003). Rather, the argument is that because there is a series of links in a causal chain leading to perceptual experience, the object in your experience is not the external object itself. One strategy to deny the significance of the above argument is to accept (1 - 3) and concede (4) but maintain that the object in your perceptual field is nonetheless directly linked to the external object by a 'causal likeness' (Descartes)[16] or a 'resemblance' (Locke). The idea here is that the external object itself produces physical disturbances which resemble its intrinsic properties at every link in the causal chain "by some parts of our Bodies, to the Brains, the seat of Sensation, there to produce in our Minds the particular Ideas we have of them" (Locke, Essay, II.viii.12). While this does not establish identity between the external object itself and the object in your perceptual field, it attempts to provide a physical, mechanistic continuity between the intrinsic properties of the object itself and the qualities of the object in your perceptual field. However, while the 'resemblance' strategy could work for *primary qualities,* as these are considered legitimate properties of matter, it is unable to deal with *secondary qualities*, as these are not considered intrinsic properties of the physical world. Locke, echoing Descartes, states the problem "[T]he ideas of primary qualities of bodies are resemblances of them, and their patterns really do exist in the bodies themselves, but the ideas produced in us by these secondary qualities have no resemblance of them at all" (E II.viii.15). Thus, even if primary qualities are somehow preserved in the causal chain from the external object to

[16] "… as a result of sensory stimulation we have a clear and distinct perception of, some kind of matter, which is extended in length, breadth and depth, and has various differently shaped and variously moving parts. . . and we appear to see clearly that the idea of it comes to us from things located outside ourselves, which it wholly resembles" (Descartes, Principles II.1, CSM i: 223).

the object in your sensory field, this does not account for the fact that the object in your sensory field presents secondary qualities. Therefore, the 'resemblance' account—known as the 'representative theory' or 'representative realism'—is of little use in thwarting the argument from causal chains because it falls prey to the argument from secondary qualities.

The causal chain of sensory processing can be leveraged in a slightly different causal argument against direct realism. In order for a causal chain to deliver the real intrinsic properties of the external object itself—that is, qualities of the object that exist independent of a perceiving subject—the causal mechanisms must maintain uniform and consistent functioning. We tend to take it for granted that our sensory physiology functions 'normally'. That is, until it doesn't function normally. Hume suggests a simple experiment. "When we press one eye with a finger, we immediately perceive all the objects to become double, and one half of them to be removed from their common and natural position" (Hume, Treatise, II). Consequently, Hume finds it obvious "that all our perceptions are dependent on our organs, and the disposition of our nerves" and takes this as a strong reason to reject direct realism (Hume, Treatise, II).

Hume's general point can be demonstrated rather drastically in terms of neuropharmacology: the chain of neurophysiological processing supporting your sensory processing rests on the neuromodulator molecules produced by your brain. You can ingest a psychedelic drug (e.g., mescaline) and the objects in your perceptual field will take on new (and sometimes very peculiar) qualities because your neuropharmacology has been altered. The direct realist is thus forced to concede that in order for you to enjoy direct perceptual contact with real properties of external objects 'as they really are' your brain must possess a specific set of 'standard' neurochemical processes that are functioning normally. This leads to problems for the direct realist: How can we precisely define the 'ordinary' neurochemical setup and its 'normal' functioning without making arbitrary assumptions? What makes us so certain that our 'normal' brain chemistry is the optimal instrument for disclosing the real properties of external objects? Endogenous neurotransmitter molecules could have evolved differently, e.g., DMT could have

become the dominant brain molecule instead of serotonin (Gallimore, 2013). Furthermore, is our concept of 'normal' perceptual faculties culturally relative? Philosopher and neuroscientist John Smythies (1956, pp. 94–5) argues that our concept of 'normal' perceptual faculties is arbitrary. "The decision to call only ordinary sense-experience real is a local phenomenon of Western European culture. It is also contingent upon the biochemical accident that our adrenal glands happen to produce adrenaline and not adrenochrome or mescaline. A valid philosophical criteria of the real should not be contingent upon cultural factors and biochemical accidents."

In summary, the argument from causal chains has two variations. The first simply states that if the external object is on one end of the causal chain of sensory processing, while the subjective experience is on the other end, then they cannot be the same entity because entities on two ends of a causal chain cannot be identical. The second variation points out the implicit assumption that our sensory causal chain must be 'functioning normally' in order to produce perceptions of objects as they actually are, and definitions of 'functioning normally' are inescapably arbitrary. The first argument from causal chains has been expressed by empirical science-minded philosophers throughout modern philosophy of perception from Descartes and Locke to Broad (1952). The second argument is much less common but no less forceful.

How does *S*-hallucination theory frame the argument from causal chains? First, I accept the argument that the concept of causal chains with mind-independent objects at one end and perceptual experience at the other end pose grave challenges for direct realism. But what should *S*-hallucination say about causality in *S*-hallucinations? In order to remain consistent, I must reject the notion that mind-independent objects cause the mind to produce the objects of perception. I must claim that mind causes (generates) the objects of perception. But what causes the mind to cause the objects? I argue that I do not need to say anything about causation. I offer a way to demarcate the veridical from the non-veridical, and that is enough. I don't start with the premise that mind-independent objects initiate a causal chain culminating in mind-dependent experience. So I don't have to say anything about what causes the mind to produce the objects of perception.

The argument from sensory latency (time-lag)

Look up at the starry night sky and it is likely that a few of the 'stars' in your visual field are in fact merely the light from stars that have actually died long ago. In the time it has taken for their light to reach your eye, they have ceased to exist, at least in the state that produced the light you see. In this case we surely can't claim that the object in your subjective visual field *is* the external object, because the object no longer exists. This is known as the time-lag argument (Broad, 1952; Fish, 2010, p. 15; Robinson, 1994, p. 80; Russell, 1948, p. 204). The idea is to generalize the burned-out-star case to cover *all* cases of visual perception: photons (and sound waves and chemicals) *always* travel at finite speeds, so there is *always* some latency before they reach your sensory organs. During this time-lag of travel latency, the object might have undergone subtle or even drastic changes. Therefore, time-lag considerations challenge the direct-realist claim that the object in perceptual experience *is* the external mind-independent object.[17]

The argument from science

The argument from science subsumes the arguments from causal chains and sensory latency, but adds many other cases where the scientific facts about the physical universe directly contradict how we experience the physical universe as it is presented by our senses. Locke (1690/2012) expresses the worry that the newly-invented microscope reveals that the physical world is very different than how the naked senses present it to us. Eddington reflects along similar lines using the much more drastic facts of contemporary physics, including the fact that his wooden table "is nearly all empty

[17] It should be noted that time-latency arguments can also be combined with facts about the functioning of neural sensory systems to produce premises that lead to similar conclusions.

space" which directly contradicts the subjective experience of solid tables (Eddington, 1935).

The argument from science has been avoided using a variety of strategies. One is to argue, for example, that a table is a 'paradigm case' of solidity and so to claim that it lacks solidity is impossible and meaningless, regardless of scientific facts (Robinson, 1994). Another is to claim that the scientific description and the properties presented in your perceptual field are each representing different—yet compatible—features of the world. A related strategy is to claim that science and perception each reveal the same world from different *perspectives* (Dummett, 1979; P. F. Strawson, 1979).

The basic premise of the argument from science is that the objects presented in our subjective perceptual field appear to have properties that science insists they do not in fact have (i.e., Eddington's table), and vice versa where science tells us that the objects have properties that we cannot perceive with our naked senses (i.e., Locke's microscope). This conflicts direct realism, which claims that the properties presented in our subjective perceptual field are the properties of the objects, independent of perceiving minds (except for secondary qualities of course).

In proposing *S*-hallucination theory I must be careful about leveraging the argument from science. Current science (mostly) assumes the existence of mind-independent objects in its descriptions of the physical universe (with the exception of certain interpretations of quantum physics, which I leave aside). *S*-hallucination theory, to be clear, only makes claims about the ontological nature of *the objects in perception* and is not (yet) making statements about the ultimate ontology of all things in the universe. However, I will have more to say about the argument from science, as it bleeds into the argument from illusion when the science of psychophysics tells us that the physical properties of certain perceptual stimuli are drastically different than the subjective percepts which they produce. In other words, *illusions*.

The argument from illusion

The argument from illusion rejects direct realism in the following way:

1. Certain physical stimuli can perceptually appear to you to bear at least some properties, *P*, which they do not actually possess (illusions[18]).
2. Apply the phenomenal principle to (1): "If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality" (Robinson 1994, p.32).
3. Since the physical stimuli does not in fact bear *P*, then the constituent of *P* in your perceptual field is *not* the physical stimuli.

At this point the argument has not 'spread' to target all cases of non-illusory 'veridical' perception and can be seen as almost tautological or merely stating the definition of illusion. The challenge to direct realism arrives via what is known as the 'spreading step' (P. Snowdon, 1992) found in (4) below:

4. We should give the same account of the constituent of your perceptual field for both veridical and illusory cases.
   Therefore,
5. In cases of veridical perception the constituent of your perceptual field is not the physical stimuli.
   Therefore,
6. direct realism is false.

Illusory perceptual phenomena are easy to come by and were discussed by Descartes, Locke, Berkeley, and Hume; however, the argument from illusion as stated above did not take center focus until the debates of early twentieth-century philosophers

[18] For a collection of illusions assembled by philosophers, see (Macpherson & Wilson, n.d.).

Broad, Moore, Price, Ayer, and others (see Robinson, 1994, pp. 31 – 58 for a detailed discussion). We will see several species of attempts to address the argument from illusion in the exposition of theories of perception outlined in Chapter 4. For now, we simply note that the argument as stated above vitally depends on (2) the phenomenal principle and (4) the spreading step.

With *S*-hallucination theory I advocate abandoning the ontological distinction between veridical perceptions and illusions. By this account, an illusion is a perceptual experience that has a lower *S*-factor in relation to a rival experience with a higher *S*-factor. Usually the perceptual experiences we call 'illusory' get labeled as such only because some other experience with a higher *S*-factor contradicts them. For example, consider the bent paddle: we say that its bentness is an illusion. Why do we say this? Because in all other circumstances the paddle appears straight; it is only when submerged in water in certain lighting conditions that it appears bent. The bent appearance has a lower *S*-factor than other rival experiences with the same content (the paddle).

While illusion is defined by cases where the physical stimuli does not bear some of the properties you experience when you encounter it, the more extreme case is when your perceptual field presents one or more objects or entities which are not in fact 'there' at all—*hallucination*.

The argument from hallucination

We have now arrived at the phenomena termed 'hallucination' and are in a position to dissect how it presents a challenge to theories of perception. According to the colloquial definition of hallucination (which isn't straightforward; see Blom, 2010), you are hallucinating when your perceptual field presents objects that are not actually there. With this definition, the argument from hallucination proceeds as follows:

1. You experience a hallucination when your perceptual field presents objects that are not actually there.

   Therefore,
2. The constituents of the hallucinatory objects presented in your perceptual field cannot be external objects because there are no external objects.
3. Both the veridical and hallucinatory cases should be given the same account of what constitutes the objects in your perceptual field.

   Therefore,
4. Direct realism is false.

We should note a few points about the argument from hallucination as stated above. First, step (1) is the colloquial definition of hallucination and step (2) simply falls out of this definition. Note that the phenomenal principle is not explicitly evoked this time because (2) implicitly assumes that there *is* some*thing* of which you are aware in a hallucination.[19] Step (3) is the 'spreading step' on which the argument hinges, which has made it the target of many counterarguments (see section Disjunctivism below). However, few can deny that step (3) is perfectly reasonable and desirable: if the subjective phenomenology you experience when you have hallucinations presenting objects which are not there is *ex hypothesi* indistinguishable from the subjective phenomenology you experience when you have veridical perceptions presenting the same objects[20], then justification for why they should require wholly distinct accounts requires defense. A further and equally reasonable assumption is that two experiences which are indistinguishable must share a 'common factor'—they have a mental state or event in common. This has been termed the 'common factor principle' summarized by the claim "Phenomenologically indiscriminable perceptions, hallucinations, and illusions have an underlying mental state in common" (Fish, 2010, p. 4). If we accept the common factor

[19] Various challenges to this aspect of (2) appear in defences of direct/naive realism (Byrne & Logue, 2008; Macpherson & Haddock, 2011; Searle, 2015)

[20] See Williamson (1990), Martin (2006), and Fish (2009, 2010) for comprehensive coverage of the intricacies of indistinguishability claims.

principle, then the spreading step (3) seems justified and the argument from hallucination holds.

A different species of argument which involves hallucinations, the common factor principle, and indiscriminibility claims is a kind of hybrid between the classic argument from hallucination and arguments from causal chains. The argument is that we could artificially stimulate your brain in the exact way that a tomato stimulates your brain and you would see a (hallucinatory) tomato. Put metaphorically in modern hacker terms, we perform a "man-in-the-middle" attack (Conti, Dragoni, & Lesyk, 2016) injecting a tomato experience at some point along the visual information pathway of physiological brain processes. From this it is concluded that external objects are not *necessary* constituents of objects presented in your perceptual field and therefore direct realism is false. This problem, articulated by Robinson (1994, pp. 153–4) and dubbed the 'screening off problem' by Martin (2004) is worth restating. First, we introduce the commonly accepted story about veridical perception which states that an external object produces a disturbance in the environment which impacts your sense organs to initiate a chain of activity terminating somewhere in your brain to present an object in your subjective perceptual field. Next we point out that this is a causal chain of activity and, per the metaphysics of causation, we could in theory butt-in at any point along the chain, stimulate the right parts, and the activity would cascade, terminate in your brain, and produce the same perception *without the stimulation of an external object*. In other words, we could make you hallucinate. The 'screening off problem' is that our causal explanation of the hallucinatory object in your perceptual field can *also* explain the veridical perception of the same object. The account of the hallucination refutes the claim that external objects a necessary component in the veridical case and so 'screens off' the need for further explanation of the veridical case over and above the explanation given for the artificially induced hallucination. Yet here the hallucination is not doing the work of the argument, thus, other than the fact that this argument involves a hypothetical hallucination, it is probably best classified as an argument from causal chains.

In summary, the classic argument from hallucination claims that hallucination presents a real problem for any theory of perception. This is because hallucinations are *ex hypothesi* indiscriminable from veridical perceptions, share a 'common factor' mental state with veridical perceptions, could in principle be induced by jumping in at any point along the causal chain assumed to support veridical perceptions, and, most agree, should for all intents and purposes be explained by the same account used to explain veridical perceptions. Yet hallucinations do not require external objects, of course; quite the opposite—hallucinatory objects *must* occur in the absence of the external objects they present.

My *S*-hallucination theory is constructed explicitly as a solution to the argument from hallucination. The most important point is that I offer a way to distinguish a veridical perception from a hallucination; that is, the *S*-factor scale. What we typically call veridical perceptions have a high *S*-factor: they are stable, successful, and in keeping with a status-quo of perceptual coherence that spans the sensory modalities of the individual and the collective of other perceiving beings. What we typically call hallucinations have a low *S*-factor: they do not evidence the same stability, successfulness, or status-quo of perceptual coherence across the collective of other perceiving beings. The two are demarcated according to their phenomenal properties and, importantly, *not* by their relation to (or lack thereof) external mind-independent objects. Thus, the argument from hallucination is not a problem for my account of what perception is.

Summary

The arguments against direct realism are many by name, but these arguments share a common format. The underlying theme is that our perceptual experiences present objects as bearing certain properties, but often these properties conflict with other perceptual experiences. This basic fact of perceptual phenomenology puts endless pressure on direct

realism, and the problem in every case is essentially a problem with the claim that external objects can possess fixed 'true' properties independent of any perceiving mind. By rejecting this claim, my account both rejects direct realism and avoids the burden of accounting for a relationship between the objects presented in perceptual experience and mind-independent objects. Furthermore, the *S*-factor scale offers hope of establishing veridical against illusory and hallucinatory perception. However, there are rival theories which both reject direct realism and attempt to account for the relationship between veridical perception and mind-independent objects. The next chapter considers these theories.

Chapter 4

Disjunctivism, intentionalism, and sense-data

The major accounts in contemporary philosophy of perception include what have been labeled sense-data theory, intentionalistic theory, and disjunctive theory. In this section I briefly state the core features of each of these theories and compare them to my *S*-hallucination theory.

Objects of perception and 'ideas'

First, it is appropriate to note the common practice of defining the 'objects' of perceptual experiences. Searle says, "In the hallucinatory state you did not see a material object, but you did see something. We need a name for these 'somethings.' In the work of Descartes, Locke, and Berkeley, they are called 'ideas.' In Hume they are called 'impressions.' In the twentieth century they came to be called 'sense data'" (Searle, 2015, p. 21). Positing ontological entities such as perceptual ideas, impressions, or sense-data is required of anyone who accepts the phenomenal principle. Why? If you uphold the phenomenal principle[21] and you try to explain the sensible qualities presented in illusion or hallucination then you need to describe the "something of which the subject is aware

[21] "If there sensibly appears to a subject to be something which possesses a particular sensible quality then there is something of which the subject is aware which does possess that sensible quality" (Robinson 1994, p.32)

which does possess that sensible quality" and in such cases you can't say it is an external mind-independent object. Descartes, Locke, Berkeley, and Hume presumably accepted the phenomenal principle (Robinson, 1994, p. 36) and were also keenly aware of the problems that perspectival variation, secondary qualities, and materialist science posed for direct realism. These philosophers found it necessary to invoke ontological entities known as 'ideas' (Locke) or 'impressions' (Hume) which have in the twentieth century become known as 'sense-data'.[22]

Sense-data theory

Sense-data theory handles the arguments against direct realism by accepting them.[23] The sense-data theorist can accept the phenomenal principle, the common factor principle, the spreading step, and the screening off problem. How? They posit an ontological entity—the sense-datum—which constitutes the objects presented in your perceptual field no matter if you are enjoying an illusion, a hallucination, or a veridical perception. You see sense-data when you see anything, no matter if what you see is veridical, illusory, or hallucinatory. The sense-data theorist does not need separate accounts for what occurs in these three cases. Sense-data are in this sense philosophically powerful because they equip the sense-data theorist with a universal mechanism which can putatively handle all of these classes of perceptual experience. However, with this explanatory power comes a great burden of responsibility—the sense-data theorist has to provide an ontological description of what a *sense-datum* actually is. Are sense-data mental or physical or neither mental nor physical? Are visual sense-data two or three-dimensional? Is subjective knowledge of sense-data fallible? What is the relationship of sense-data to concept-involving activities such as recognition? Are sense-data particulars or

[22] Searle (Searle, 2015) calls this move "One of the biggest mistakes in philosophy in the past several centuries."
[23] In fact, most of the arguments against direct realism have been articulated by sense-data theorists to motivate the theory.

universals? Are sense-data modes of sensory activity or objects of it? Most importantly, sense-data theory does not deny the existence of mind-independent objects, so it has to account for the relationship between sense-data and the external mind-independent object that it 'resembles'.

Realism in sense-data theories

In addition to the burden of giving an account of the ontological status of sense-data, the sense-data theorist also carries the burden of providing an account of the *relation* between a subjective, mind-dependent sense-datum and an objective, mind-independent external object. Locke held that sense-data (sensory 'ideas') are *caused by* the external objects and thus inherit some of their properties so that the sense-data bear a relation of *resemblance* or *representation* to the external physical world. Locke's account, to some extent also held by Descartes, is termed *representative realism*.

Phenomenalism and idealism in sense-data theories

By contrast, Berkeley, Hume, and Mill oppose representative realism. Instead, they argue that what we call the 'external physical world' is no more than the subjective, mental, internal sense-data that arise in beings capable of having perceptual experiences. This view has come to be termed phenomenalism or idealism. The phenomenalists agree with the representational realists that sense-data are the constituents of the objects in your sensory field, but disagree about their relation to the external physical world. "According to the idealists and phenomenalists, such as Berkeley, Hume, and Mill, the data are actually part of physical objects, for objects consist only of actual or actual and possible sense-data …" (Robinson, 1994, p. 3).

My *S*-hallucination theory sides with Berkeley, Hume, and Mill and lands on the phenomenalist side of the realism-phenomenalism divide. Furthermore, I adopt a Kantian-like account of the origin of sense-data in that I claim that sense-data are produced in the mind, generated by the brain.

For various reasons which I will not get into here, sense-data theories fell out of favor in the second half of the twentieth century. In the late 1960s, two rival theories emerged simultaneously and continue to dominate debates in philosophy of perception; namely, *intentionalistic* theory (Anscombe, 1965) and *disjunctive* theory (Hinton, 1967a, 1967b). I cover each respectively in the sections that follow.

Intentionalistic theory

Intentionalistic theory is a broad term which applies to any account of perception holding that your perceptual field is made up of properties which are *about* external objects and states of affairs. The intentionalistic theorist generally holds that perceptions are analogous to beliefs or propositional attitudes in that they have *intentional* content that represents[24] the world as being a certain way. A handy aspect of intentionality is that an intentional object can be 'about' something that doesn't actually exist—not all beliefs are true beliefs. The intentionalistic theorist thus has a device which can consistently handle questions about what constitutes the objects in your perceptual field no matter if they are veridical, illusory, or hallucinatory. In the veridical case, the object in your visual field is *true*; in the illusory case it is *partially true*, and in the hallucinatory case it is *false*. The intentionalistic theorist can thus promote a clean mapping that appeals to common-sense: veridical perceptions *represent*, illusions *mis*represent, and hallucinations represent something that does not exist, similar to the way in which thinking of a unicorn can represent a unicorn but does not require the existence of a unicorn. According to the intentionalistic theorist, then, perceptual experiences have 'correctness conditions' or 'accuracy conditions' or 'conditions of satisfaction' where these features *do not require that the condition be fulfilled*. Note that this aspect of intentionalism can appear to circumvent the phenomenal principle. That is, the intentionalist hacks the phenomenal principle to read more like this: 'If there sensibly appears to a subject to be something

[24] Intentionalistic theories of perception talk of perceptions as having representational properties but should not be confused with the 'Representative Theory' of perception or 'Representative Realism' which hold that *sense-data* represent or resemble states of the world (see section Sense-data theory above).

which possesses a particular sensible quality then there is something of which the subject is aware which ~~does possess~~ *represents* that sensible quality'. This renders the phenomenal principle much less metaphysically demanding because the subject does not need to be aware of something that *has* the properties; it merely *represents them*. Importantly, recall that the majority of the arguments in the previous section turn on accepting the phenomenal principle. Thus the intentionalistic theorist appears to have a way to escape the threat of the argument from illusion, hallucination, perspectival variation, etc. by rejecting—or neutralizing—the phenomenal principle.

Furthermore, the intentionalistic theorist can embrace the common-factor principle *and* the 'spreading step' by espousing a common factor—aboutness or intentionality or representation—which constitutes the objects in your perceptual field across veridical, illusory, and hallucinatory cases.

Intentionalistic theories of perception came to dominate philosophy of perception beginning in the late 1970s (Crane, 1992) and have steadily increased right up to the present day (Brogaard, 2014). The volume of discussion around intentionalistic theories contains intricacies which lie outside of our scope here. In summary, a large portion of these discussions center around the questions of how *concepts* feature in an intentional account of perception; what the *content* of a perception actually is; whether the content is 'conceptual' or 'nonconceptual'; and the relationship between the subjective phenomenology of perception and its intentional content.

One of the central objections to the intentionalistic theory is that it does not offer a demarcation between perceptual experience and more familiar forms of mental representation such as belief, thought, or even mental imagery. Why would the intentionalistic theorist need to demarcate the intentionality of perception from the intentionality of thought? One reason is the differences in subjective phenomenology. Imagining a unicorn is not the same as seeing one trotting past you on a majestic green hillside. Thinking about a pineapple is not the same as tasting one. The intentionalistic theorist thus has a responsibility to explain why you experience the contents of your perceptual field as direct presentation of actual objects and not as being 'about' external

objects. However, thus formulated, this objection applies to all theories that are not direct realism, including sense-data theory. A related phenomenological objection which applies uniquely to intentionalistic theory can be raised by pointing out that that objects in your perceptual field are presented with a definitiveness that is not shared with other intentional objects in your mind, such as your thoughts, beliefs, or mental imagery. The intentionalistic theorist has to explain why perceptual content is not as vague and pliable as these other forms of intentionality. One response might be that the intentional powers of perceptual experience have the unique ability to produce 'transparent' contents that *represent* objects *as* 'here' 'present' 'concrete' as part of their conditions of satisfaction. Again, this is handy because it can appear to avoid accepting the phenomenal principle and thus avoid the arguments which turn on that principle. However, the intentionalistic theorist is now required to articulate the metaphysics of these intentional perceptual objects—their ontological status—in such a way that distinguishes an 'intentional object' from a sense-datum. Furthermore, there is the question of how perceptual contents perform this neat trick not found in other intentional objects. "It would be fair to say that the whole problem for 'intentional object' theories of perception is whether sense can be made of the idea of an entity that can do this; that is, represent its content transparently … " (Robinson, 1994, p. 8).

Realism in intentionalistic theories

Intentionalistic theories describe perceptual experiences as being of or about some state of affairs in the world—the objects in your perceptual field are individuated by their intentionality. To establish a mapping to the direct realism of pre-philosophical intuitions about perception, the intentionalistic theorist must show how the intentionality of perceptual experience is linked to the actual states of the world presented in the experience. Many philosophers accept the charge that the intentionalistic approach contradicts direct realism because it introduces a medium—a 'veil of perception'—between the perceiver and the world. It is common (and reasonable) to assume that when intentionalistic theorists champion the common-factor principle

(especially when responding to the argument from hallucination) they are implicitly rejecting direct realism. Furthermore, when the intentionalistic theorist advocates 'intentional content' as the common factor shared between veridical, illusory, and hallucinatory cases, their explanation of the hallucinatory case 'screens off' the possibility of an explanation on which external objects are fundamental to the veridical case. Taken together, these points are widely cited as reasons to classify the intentionalist theories as forever divorced from direct realism.

However, Searle (2015) has recently argued for an intentionalistic theory of perception that explicitly promotes direct realism. Searle defends direct realism, rejects the phenomenal principle, rejects the argument from hallucination, yet accepts the common factor principle. How does Searle address the argument from hallucination? Searle's key strategy is to "expose" what he calls "the fallacy" which lies at the heart of what he labels "The Bad Argument" which Searle describes as "One of the biggest mistakes in philosophy in the past several centuries." What is The Bad Argument? None other than the argument from illusion, the argument from hallucination, or any of the arguments which take that general form—so, most of the arguments in Chapter 3 above. And what is 'the fallacy' upon which "The Bad Argument" supposedly rests? None other than the phenomenal principle! What is fallacious about the phenomenal principle? Searle argues that it makes a critical mistake by confusing the object of awareness with the awareness itself. Here is how I think Searle would rewrite Robinson's description of the phenomenal principle: 'If there sensibly appears to a subject to be something which possesses a particular sensible quality then ~~there is something of which the subject is aware which does possess that sensible quality~~ *the subject is having a perceptual experience which presents objects and states of affairs in the world as having that sensible quality.*' Searle's point is that when you render the situation in this terminology—when you remove the fallacy resulting from the confusion between perceptual experience and the objects of the experience—it becomes philosophically unproblematic. Thus, Searle argues, all of the arguments against direct realism lose

traction without the fallacy. Freed of the constraints of the "Bad Argument," Searle builds an intentionalistic direct realist account of perception.

Like all intentionalistic accounts of perception, Searle's account uses intentionality to conveniently explain what happens in hallucination. Searle simply holds that in every case of perceptual experience (veridical or otherwise) there is an ontologically subjective experience, caused at least in part by the powers of intentionality possessed by brains. In the veridical cases, an essential causal element of the subjective experience resides in the external mind-independent 'ontologically objective' object. In the non-veridical cases, the subjective experience is caused entirely by internal brain processes of intentionality and no external object is required.

My offering of *S*-hallucination theory has some things in common with Searle's intentionalistic offering and with other intentionalistic theories in general. First, I hold that the brain brings intentionality and meaning to every experience. Second, I hold that hallucination and illusion are the result of these intentionality-generating powers of brains. However, I depart from Searle and other intentionalistic theories in that I reject the claim that the constituents and defining features of veridical perceptions are mind-independent objects. The definition of veridical perception, I argue, must be defined in terms of *S*-factor characteristics. This approach makes my theory resistant to attacks based on the common-factor principle leading to the screening-off problem. Recall that the screening-off problem is a label for the situation in which a theory of perception accounts for hallucination and illusion in such a way that the need for external objects is 'screened off'—whatever the theory invokes to explain hallucination should also explain veridical perception. Intentionalistic theories, I conclude, are too susceptible to the screening-off problem. As an illustration, compare Searle's Figure 1.2 with Figure 1.3. Why is the external object shown in Figure 1.2 necessary when the "Internal causes" and "Intentionality" of Figure 1.3 are sufficient to produce visual experience in hallucination?

The force of the screening-off problem, combined with the unrelenting desire to uphold direct realism, is what motivates disjunctive theory.

Disjunctive theory

Recall Hume's sentiment that "this universal and primary opinion of all men is soon destroyed by the slightest philosophy" (Hume, 1777, p. 151) or Reid's remark that rejecting naive realism appeared to him to be "a unanimity … among philosophers" (Reid, 1785). Disjunctive theory breaks this trend[25] and provides an account designed to *uphold* the ordinary conception. How? How can one maintain direct realism against the power of the perennial arguments from causal chains, illusion, and hallucination described in the previous chapter? The disjunctive theorist sticks to their guns with the following strategy. They deny that there is a common factor shared by your veridical, illusory, and hallucinatory experiences. Your perceptual field is either veridical or hallucinatory; one or the other; the choice of ontological categories is *disjunctive*. You look at your countertop and have a visual experience of a bowl of lemons. *Either* you see a bowl of lemons, *or* you hallucinate a bowl of lemons. The former experience is of the veridical perception kind and requires mind-independent lemons; the latter experience is of the hallucination kind and requires the *absence* of a mind-independent lemons. Crucially, says the disjunctive theorist, the veridical perception and the hallucination "*do not* have an underlying mental state in common" (Fish, 2010, p. 89).

So disjunctive theory at its core is defined by a rejection of the common factor principle. It denies common elements between veridical perceptions and hallucinations. How does the disjunctive theorist know that you are having a veridical perception? They call it veridical when the constituents of your experience are actual external mind-independent objects. How does the disjunctivist know when you are hallucinating? They call it hallucination when the constituents of your experience are *not* external mind-independent objects. Disjunctive theorists argue that because they defend naive

[25] Disjunctive theories are relatively recent in philosophy of perception, marked primarily by the work of Hinton (1967a, 1967b), Snowdon (1980), McDowell (1983, 1986), and Martin (2006; 2002, 2004).

realism, their account is the 'least revisionary' philosophical account compared with the ordinary conception.

What are the consequences of the disjunctive position? First, the disjunctive theorist has difficulty saying anything at all about illusion or hallucination. These are defined negatively, as being not completely veridical perceptions. In the case of hallucination, this might squeak by (though see below). In the case of illusion, however, the disjunctivist account has little to say about why it is that certain physical stimuli are presented in your perceptual field as having properties which they do not have. "Naive realists might be thought to be making either or both of two mistakes. They might be thought to be failing to notice that objects often do look other than they are: or they might be accused of not realising that perceiving objects directly should entail seeing them as they are. The main point is, I think, the latter" (Robinson, 1994, p. 41).

Second, the disjunctive theorist faces the challenge of denying the common-factor principle without destroying the very definition of hallucination. Hallucination is almost unavoidably defined as a perceptual experience, indistinguishable from, or at least similar to, a veridical perception, which occurs in the absence of external physical stimuli. Many disjunctive theorists, however, hold that there is *nothing* in common between your veridical perception and your hallucination *even if you can't distinguish between the two*.

Finally, the disjunctivist becomes forced to deny what seem to be uncontroversial overlaps between hallucination, illusion, and perception. For example, since a veridical perception has subjective phenomenal character, some disjunctivists have realized that this commits them to denying that hallucinations have any phenomenal character—*at all* (Fish, 2009). This trend is not limited to phenomenal properties. The intentionalistic theorist, for example, posits that there is at the very least an overlap in intentional content between the veridical and the hallucinatory cases. This is part of the appeal of intentionalistic theories. The disjunctive theorist, on the other hand, must deny that a veridical perception and its corresponding hallucination can share intentional content. Unsurprisingly, many thinkers find this position absurd. "If you deny that the hallucinatory experience and the veridical experience—the good case and the bad

case—can be phenomenologically and intentionalistically exactly the same, then you know you have made a mistake. If you have a theory that has that as a consequence, then you know the theory is false because it entails a false proposition. Amazingly, there is a class of philosophers of perception who accept precisely this *reductio ad absurdum*. They are called Disjunctivists" (Searle, 2015, p. 164).

Like Searle, I cannot tolerate the absurd consequences of disjunctive theories. However, my *S*-hallucination theory has something in common with disjunctivism in that I accept the force of the screening-off problem. Disjunctivists handle this by rejecting the common-factor principle in order to uphold an account of mind-independent objects, direct realism. I take the opposite approach: I reject mind-independent objects in order to uphold the common-factor principle *and* avoid the screening-off problem.

Summary

Sense-data theories come in realist or phenomenalist variants. The former tries to account for how sense-data relate to mind-independent objects, while the latter argue that sensible objects only make sense in the context of a perceiving mind. My theory has much in common with phenomenalist sense-data theories. Intentionalistic theories hold that the proper way to conceive of the relationship between the objects of perception and mind-independent objects is that perceptual experience is *about* mind-independent objects; it holds them as its conditions of satisfaction. Intentionality is a power of minds and brains and thus can account for hallucination. Moreover, intentionality is claimed to be the 'common factor' between all forms of perceptual experiences, veridical or otherwise. My theory is sympathetic to talk of the intentionalistic powers of minds and brains, and to this being a common-factor between veridical and non-veridical perception, but rejects (Searle's) intentionalist claims that veridical perception is defined by the conditions of satisfaction being satisfied by mind-independent objects. Disjunctive

theories cite the screening-off problem as a deal-killer for intentionalistic theory. My theory is in agreement with disjunctivism on this point. However, disjunctivists reject the common-factor thesis in hopes of upholding the direct realist claim that mind-independent objects are required constituents of true veridical perception. Because this move leads to absurd consequences like denying that hallucinations have phenomenal character, I take exactly the opposite approach: I reject the very idea of mind-independent objects and this allows me to uphold a strong version of the common-factor principle. My *S*-factor scale is novel in the context of this debate because I offer a theory that is able to demarcate veridical from non-veridical perceptions, uphold the common-factor principle, and avoid the screening-off problem. In the next chapter I provide a more detailed sketch of how *S*-hallucination theory accomplishes this.

Chapter 5

Making sense of reality

In Chapter 2 I claimed that all of the objects presented in your perceptual field, veridical or not, are always produced in your mind, generated by your brain, and thus cannot be considered 'mind-independent' in any meaningful way. Furthermore, I made the negative claim that mind-independent objects should not ground our criterion for distinguishing veridical from non-veridical experiences. To fill the gap left by denying us of the traditional mind-independent criterion, I offered a metric, which I call the *S*-factor scale, for establishing how much veridicality a given experience can be said to exhibit. *S*-factor is determined by analyzing an experience for three key '*S*-traits' which are ontologically subjective and thoroughly experiential: stability, success, and status-quo. A given perceptual experience is maximally veridical when it exhibits maximum levels of all three of these traits. Veridicality is reduced as any of these traits diminishes. This, I claim, is the only type of criterion possible for establishing how much 'reality' is present in a given experience.

My claim that all objects in perceptual experience are mind-dependent and that all criteria for veridicality must also be mind-dependent is likely to raise objections to the effect that I am abandoning all hope of objective reality, reducing our perception of objects to mere 'illusions on the fancy'. My aim in this chapter is to assuage this worry and show how experiences with high degree of *S*-traits can be as secure as ever, and as real as ever.

First, I must admit that most of the perceptual experiences that we hope to call veridical actually exhibit a sort of direct realist 'feel' *as part of their subjective phenomenal character*.

Hume (Hume, 1777, sec. XII.1 118) articulates this fact:

> It seems evident, that men are carried, by a natural instinct or preposession, to repose faith in their senses; and that, without any reasoning, or even almost before the use of reason, we always suppose an external universe, which depends not on our perception, but would exist, though we and every sensible creature were absent or annihilated. Even the animal creation are governed by a like opinion, and preserve this belief in external objects, in all their thoughts, designs, and actions.
>
> It seems also evident, that, when men follow this blind and powerful instinct of nature, they always suppose the very images, presented by the senses, to be the external objects, and never entertain any suspicion, that the one are nothing but representations of the other.

In philosophy of perception this has been called 'openness to the world' (McDowell, 1994) or 'perceptual presence' which PF Strawson (1979, p. 97) describes as the fact that "mature sensible experience (in general) presents itself as … an immediate consciousness of the existence of things outside us." In other words, the subjective phenomenal character of our perceptual experience itself exhibits the 'feeling' that the objects it presents are really there and would be there independent of your perceiving them.

This opens up a neat objection to my whole line of argument. *Could we not use the S-factor scale to establish the 'reality' of direct realism?* After all, 'openness' or 'perceptual presence' is a subjective experiential quality. Moreover, it has an incredibly high degree of *S*-traits! It occurs in a repeatable manner and remains durable across a variety of perspectives (stable); it enables action or actionable information (successful), and it is consistent with past experience and with prior knowledge about 'how things are'

and ‘how the world works’ (status-quo). It is so status-quo, in fact, that, as Hume (1777, sec. XII.1 118) points out, “Even the animal creation are governed by a like opinion, and preserve this belief in external objects, in all their thoughts, designs, and actions.” Therefore, it would seem that my *S*-factor scale, when applied to the phenomenal character of ‘openness’ found in the majority of perceptual experiences, refutes the very core ontological claim I make with *S*-hallucination theory; namely, that all objects are mind-dependent. Put another way, by my own *S*-factor scale, the ubiquitous subjective experience of the mind-independence of external objects is rendered *highly veridical*!

However, this point is not a refutation of my claims, for two primary reasons. First, the phenomenal character of ‘openness’ or ‘perceptual presence’ does not establish that the objects have mind-independent ontology. As PF Strawson argues,

> The reason for this is simple. In order for some belief or set of beliefs to be correctly described as a theory in respect of certain data, it must be possible to describe the data on the basis of which the theory is held in terms which do not presuppose the acceptance of the theory on the part of those for whom the data are data. But this is just the condition we have seen not to be satisfied in the case where the so-called data are the contents of sensible experience and the so-called theory is a general realist view of the world. The “data” are laden with the “theory” (P. F. Strawson, 1979, p. 95).

So the data of perceptual experience, even though they are infused with realist phenomenal character, and even though this character qualifies for a very high *S*-factor on my scale, can in no way be leveraged as evidence in support of the theory that veridical perception is grounded in objects with a mind-independent ontology. No mind-dependent experience of an object can prove that the object possesses mind-independent ontology. And all experience is mind-dependent.

Second, the consideration that openness and presence score high *S*-factor actually renders more palatable my central argument that the ‘reality’ and veridicality of perceptual experiences are properly grounded in the *S*-traits of their subjective phenomenal character. My claim, again, is that the higher the *S*-factor of an experience, the more ‘reality’ is ‘in’ it. So when we consider that most humans and most animals

experience their perceptual field with a phenomenal character that robustly exudes a realist, mind-independent objective world, then by my criterion this experience is real and veridical. Again, as pointed out by PF Strawson above, these perceptual-phenomenological data can in no way establish their own mind-independent ontology. Importantly, the high *S*-factor makes them, on my theory, more 'real' and 'veridical' than, say, the more rare (and less stable, successful, or status-quo) experience where your perceptual field is experienced as 'not real' or 'mind-dependent' or 'all in the mind'. This key point can calm the worry that without some connection to mind-independent ontology the objects presented in perceptual experience are reduced to 'mere illusions on the fancy'. Many philosophers worry that anything mind-dependent is too pliable, too willy-nilly, too variable to be properly called 'reality' or veridical. Mind-dependent entities can be willed in or out of being at your fancy, the worry goes, so this is a hopeless basis for true reality or veridicality. This worry can be soothed, I argue, by considering the very fact that openness and presence are very stable, very successful, and very status-quo features of the phenomenal character of perceptual experience. Thus, even if we don't get to ground them in objective mind-independence, we still get to say they are more 'real' than illusions on the fancy. So the ubiquity and high *S*-factor of openness and perceptual presence—the subjective feeling that *those* objects are part of an objective mind-independent world, which occurs in a wide variety of perceiving beings—allows us to say that the contents of these experiences are as secure as ever and as real as ever.

But, you might object that my definition of 'veridical' is in trouble here. You might charge that if I claim that openness and presence and all of their mind-independent-suggesting character are rendered veridical on my theory, yet the theory also claims that objects with mind-independent ontology are not possible in perceptual experience, then I would be saying that the experience is veridical *and* false at the same time. My reply is that I have no problem with an experience being perceptually veridical yet philosophically false.

In summary, my view is that billions of perceiving beings experience objects that are in fact mind-dependent but which have the phenomenal character of being robustly mind-independent. Because this experience is so ubiquitous, stable, successful, and status-quo, it counts as veridical and real on my theory. At the same time, any mind-*independent* ontology derived solely from some aspect of phenomenal experience is metaphysically false on my theory.

_S_-hallucination theory in action

As mentioned in the Introduction, often our theories of perception are driven by cases of non-veridical perceptual experience. As with other theories, a productive way to illustrate *S*-hallucination theory is to use non-veridical cases of perceptual experience. In this section I use examples of non-veridical or partially-veridical perceptual experience to illustrate the strengths of the *S*-hallucination theory of perception.

Illusion

Which line is longer in Figure 1?

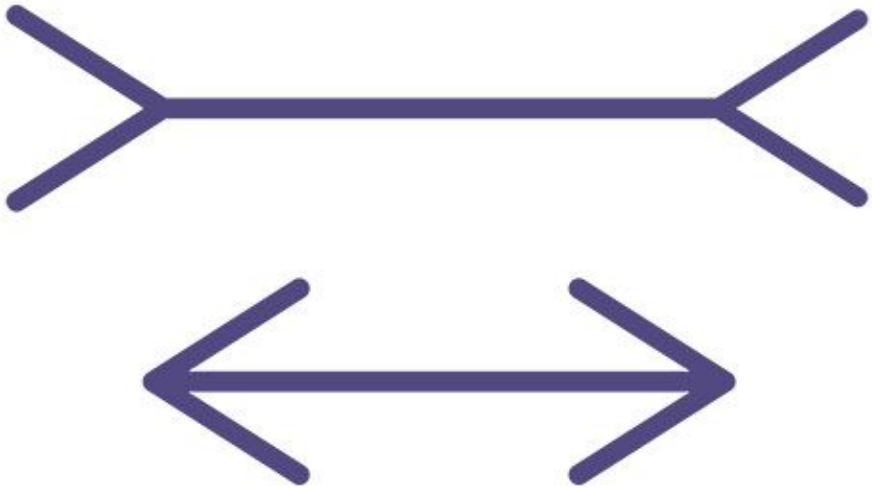

Figure 1: The Müller-Lyer illusion (Macpherson & Wilson, n.d.). See https://www.illusionsindex.org/ir/mueller-lyer

To most viewers, it seems clear and certain that the top line is longer than the bottom. However, the top line and the bottom line are exactly the same length. So in some meaningful sense your perception is not veridical. Fortunately the *S*-factor metric is a gradient scale. Your experience scores high on the *Stability* item; the top line looks longer every time. It is also high on the *Status-quo* item; the top line looks longer to most viewers (although see Segall, Campbell, & Herskovits, 1963). However, the *Successful* item will score low. If these lines were bars of gold and you were told the top bar costs more because it is longer, you could be fooled into paying more money for the top line

yet receive no extra gold for your money. Thus, your perception is not very successful. S-*factor* is neat because its gradient-scale nature allows us to say that your illusory experience of Figure 2 is 'partially veridical'—and that is exactly what we intuitively want to say about illusions. Other theories struggle to provide a characterization of illusion that is consistent with their characterization of perception and hallucination. An objection might be raised that, according to the *S*-factor scale, many well-known illusory phenomena would score a high *S*-factor. Let us run through the criteria for *S*-factor to make this point explicit. Illusions can exhibit very high levels of *Stability* across individuals and even across species. Illusions can indeed enable *Successful* actions and provide actionable information. Illusions can even be consistent with the 'status quo' of past experience and prior knowledge of how the world works. In fact, the leading explanations of optical illusions in perceptual neuroscience cite past experience, environmental utility, and prior knowledge of 'how the world is' as *causes* of the illusion (Geisler & Kersten, 2002; Gregory, 1968; Hoffman, 2012; Purves et al., 2015).

Hallucination

The classic philosophers' example of hallucination is a pink rat—or pink elephant—which appears in your perceptual field as if it is an external, mind-independent entity, but is in fact a hallucination. The additional assumption is that in this situation you are the only perceiving being who 'sees' the rat and others see the same scene sans rat. It is a puzzle as to why the rat is a mind-dependent object embedded into a broader perceptual scene of objective mind-independent objects. The puzzle is philosophically arduous for all of the reasons detailed in the arguments covered in Chapter 3. What does my theory say about the veridicality of such an experience? It is simple. *S*-factor is high for the shared experiences of the perceptual scene; the objects that all can witness, so these experiences are veridical. The *S*-factor is low for the experience of the pink rat. Thus, the pink rat is non-veridical, a hallucination. My theory gives the result we want—the ability to justify calling the pink rat hallucinatory—without invoking direct realist criteria of veridicality. The veridical scene is mind-dependent, the hallucinatory rat

is mind-dependent in the exact same way. But the rat is a hallucination and the surrounding scene is not.

Psychedelic drug effects

Hallucination in the real world is rarely like the pink rat scenario. Psychedelic drugs, which are neuromodulator molecules found in nature, cause your visual experience of the external world to take on unusual characteristics (Swanson, 2018). Colors gyrate their apparent brightness and contrast, objects shrink or grow or distort, faces morph into bizarre configurations, everyday objects morph into lively faces. All of this due to a small molecule impacting neurons in your brain. This puts pressure on traditional philosophical theories of perception, which typically write such phenomenology off as 'abnormal'. Direct realism is especially dependent on the claim that the brain must be functioning 'normally' to obtain veridical experiences of mind-independent objects. So their claims often must be prefaced with "If you have normal vision and are in reasonably good light …" (Searle, 2015, p. 11). "When vision is correctly doing its biological job …" (Searle, 2015, p. 13). But what are the correctness conditions here? Your brain runs on endogenous small molecules. You can tweak these by, for example, ingesting psilocybin, after which you will begin to experience what are usually static objects but which are now suddenly pulsating, melting, or winking at you. How can we demarcate 'normal vision' 'correctly doing its biological job' from 'abnormal vision' not doing its correct job? My *S*-hallucination theory handles this with ease. I say that your normal waking perceptions of an external scene and your trippy psilocybin perceptions of that same scene are *both* produced in your mind, generated by your brain. Psilocybin makes your mind generate unusual experiences, but your normal experiences have the same mind-dependent status. Of course, veridicality is low with the psilocybin experience because its phenomenal characteristics lack the stability, success, and status-quo of perception-as-usual. Importantly, my theory is compatible with the claim that a single scene can contain a mixture of veridical and non-veridical elements. What matters for veridicality of any given element is its *S*-traits.

Perceptual destabilization stimuli

One final example, rarely considered by philosophers, is the phenomenon of *perceptual destabilization*. "Prolonged sensory restriction and monotonous forms of stimulation, such as exposure to a luminous flickering field, can induce a range of perceptual effects. In the visual modality these can vary from a kaleidoscopic medley of patterns (e.g., spirals, honeycombs, lattices), often accompanied by illusory motion and color, all the way to more complex, dreamlike images" (Miskovic et al., 2019, p. 266). Importantly, these non-veridical perceptual experiences result from "exposure to sensory patterns that deviate substantially from the statistics of natural environments" (Miskovic et al., 2019, p. 265). This phenomenon is strongly consistent with my theory, and, I argue is evidence that your mind uses something like *S*-factor to bootstrap its production of veridical experiences. To induce destabilization, we introduce perceptual experiences with a low *S*-factor—"sensory patterns that deviate substantially from the statistics of natural environments"—they lack stability and status-quo because they conflict with statistical regularities of previous experiences. Your mind loses its ability to ground its perceptual productions in *S*-traits, so it slowly drifts from veridical to non-veridical activity. This, I argue, provides evidence that veridicality is established by *S*-traits.

Summary

We can establish criterion for veridical reality without invoking talk of mind-independent objects. I claim that every object presented in your perceptual field is mind-dependent. We can never, I argue, identify some aspect of your perceptual experience that is mind-independent. But we can identify aspects that are veridical and real. To do this, we analyze the patterns in your perceptual experience and rate them for their relative level of

S-traits. If you look to your right and see a pile of rocks and this pile is there every time you look and every time your fellow perceiving beings look; if you can touch, smell, hear, and taste the rocks; if the rocks are coherently situated within your broader expectations and understanding of the world—then your experience is veridical. If the rocks are jiggling, pulsating, floating in the air and sprouting faces and winking at you while your fellow perceiving beings do not see this, then your experience is non-veridical.

The fact that perceptual experience has the phenomenal character known as 'openness' or 'perceptual presence'—the feeling that you are experiencing an direct-realist-style objective mind-independent world—and that this character has a high *S*-factor, does not refute my theory. Rather, the ubiquity of this phenomenal character grounds the veridicality of those experiences in high *S*-traits, soothing worries that a wholly mind-dependent perceptual ontology is nothing but illusions on the fancy.

S-hallucination theory can handle the fact that your perceptual field can contain elements of mixed-veridicality: illusions, hallucinations, partial hallucinations, and psychedelic visuals.

Conclusion

In the present work I have argued that objects with a mind-independent ontology have no place in a theory of what the objects of perception *are*. I insist that our concept of hallucination is a better metaphor for the actual dynamics behind perceptual experience than our traditional realist conception of perception. I proposed that *S*-traits are the mark of veridicality, not mind-independent objects. The *S*-factor metric, I argue, is what the mind itself uses to deliver its own phenomenal character with a 'sense of reality' in a perceptual experience, and, furthermore, that this metric is our best basis of a philosophical theory of what constitutes veridicality in perceptual experience. I have chronicled the classic arguments against direct realism, the theories which have tried to account for veridicality, and the varieties of perceptual experience that these theories fail to account for in a satisfying philosophical way. The overall picture that results from my account is that the mind 'makes sense' of its own phenomenology using something like *S*-traits to demarcate 'reality' from non-veridicality, and that we as philosophers can make sense of hallucination, illusion, perception, and everything in between if we adopt a theory whereby all objects are ontologically mind-dependent but that some experiences nonetheless can be classified as 'more real' or 'more veridical' than others. All perceptual experience is hallucination-like, but those experiences with high *S*-factor *are* reality, and, even though I rather dramatically call them *S*-hallucinations, their veridicality remains as secure as ever, and as real as ever.

References

Anscombe, G. E. M. (1965). The Intentionality of Sensation: A Grammatical Feature. In R. J. Butler (Ed.), *Analytic Philosophy* (pp. 158–180). Blackwell.

Atalla, J., Orwig, J., & Salter, L. (2018, March 26). A neuroscientist explains why reality may just be a hallucination. *Business Insider*. Retrieved from https://www.businessinsider.com/neuroscientist-explains-why-reality-hallucination-meaning-2018-3

Berkeley, G. (1709). *An Essay Towards a New Theory of Vision*. Aaron Rhames.

Berkeley, G. (1710). *A Treatise on the Principles of Human Knowledge*. Aaron Rhames.

Berkeley, G. (1713). *Three Dialogues Between Hylas and Philonous*. Oxford University Press.

Blom, J. D. (2010). *A Dictionary of Hallucinations:* Springer New York.

Bornat, R., & Brady, J. M. (1976). Using knowledge in the computer interpretation of handwritten FORTRAN coding sheets. *International Journal of Man-Machine Studies*, *8*(1), 13–27.

Broad, C. D. (1952). Some Elementary Reflexions on Sense-Perception. *Philosophy* , *27*(100), 3–17.

Brogaard, B. (2014). *Does perception have content?* Oxford University Press.

Byrne, A., & Logue, H. (Eds.). (2008). *Disjunctivism*. MIT Press.

Caston, V. (2015). Perception in Ancient Greek Philosophy. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception*. Oxford University Press.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, *36*(3), 181–204.

Clark, A. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press.

Conti, M., Dragoni, N., & Lesyk, V. (2016). A Survey of Man In The Middle Attacks. *IEEE Communications Surveys Tutorials*, *18*(3), 2027–2051.

Crane, T. (1992). *The Contents of experience : essays on perception*. Cambridge: Cambridge University.

Crane, T. (2005). What Is the Problem of Perception? *Synthesis Philosophica*, *20*(2), 237–264.

Crane, T., & French, C. (2017). The Problem of Perception. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2017). Retrieved from https://plato.stanford.edu/archives/spr2017/entries/perception-problem/

Dayan, P., Hinton, G. E., Neal, R. M., & Zemel, R. S. (1995). The Helmholtz machine. *Neural Computation*, *7*(5), 889–904.

Descartes, R., Cottingham, J., Stoothoff, R., & Murdoch, D. (2008). *The philosophical writings of Descartes. Vol. 2 Vol. 2*. Retrieved from https://www.worldcat.org/title/philosophical-writings-of-descartes-vol-2/oclc/731000142&referer=brief_results

Dummett, M. (1979). *Common Sense and Physics*. 1–40.

Eddington, A. (1935). *The Nature of the Physical World*. London: Dent.

Fish, W. (2009). *Perception, hallucination, and illusion*. OUP USA.

Fish, W. (2010). *Philosophy of perception: A contemporary introduction*. Routledge.

Flores, J. K. (2015). Hume in the Light of Bayes: Towards a Unified Cognitive Science of Human Nature. *Res Cogitans*, *6*(1), 2–13.

Frith, C. D. (2007). *Making up the mind : how the brain creates our mental world*. Malden, Mass.: Blackwell.

Galilei, G. (1623). Il saggiatore. In A. C. Danto & S. Morgenbesser (Eds.), *Philosophy of Science*.

Gallimore, A. R. (2013). Building Alien Worlds—The Neuropsychological and Evolutionary

Implications of the Astonishing Psychoactive Effects of N, N-Dimethyltryptamine (DMT). *Journal of Scientific Exploration*, *27*(3), 455–503.

Geisler, W. S., & Kersten, D. (2002). [Review of *Illusions, perception and Bayes*]. *Nature neuroscience*, *5*(6), 508–510.

Glaser, D., & Sanderson, M. (2018, April 2). A Neuroscientist Explains: where perception ends and hallucination begins - podcast. *The Guardian*. Retrieved from http://www.theguardian.com/science/audio/2018/apr/02/a-neuroscientist-explains-where-perception-ends-and-hallucination-begins-podcast

Gregory, R. L. (1968). Perceptual illusions and brain models. *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character. Royal Society* , *171*(1024), 279–296.

Grush, R. (2004). The emulation theory of representation: motor control, imagery, and perception. *The Behavioral and Brain Sciences*, *27*(3), 377–396; discussion 396–442.

Hatfield, G. (2015). Perception in Philosophy and Psychology in the 19th and Early 20th Centuries. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception*. Oxford University Press.

Helmholtz, H. von. (1867). *Treatise on Physiological Optics Vol. III*. Dover Publications.

Hinton, J. M. (1967a). Experiences. *The Philosophical Quarterly*, *17*(66), 1–13.

Hinton, J. M. (1967b). Visual Experiences. *Mind; a Quarterly Review of Psychology and Philosophy*, *76*(April), 217–227.

Hoffman, D. D. (2012). The Construction of Visual Reality. In J. D. Blom & I. E. C. Sommer (Eds.), *Hallucinations* (pp. 7–15). Springer New York.

Hoffman, D. D., & Prakash, C. (2014). Objects of consciousness. *Frontiers in Psychology*, *5*, 577.

Hohwy, J. (2013). *The predictive mind*. OUP Oxford.

Hohwy, J. (2016). The Self-Evidencing Brain. *Noûs*, *50*(2), 259–285.

Horn, B. K. P. (1978). *Vision review*. Retrieved from https://dspace.mit.edu/bitstream/handle/1721.1/41985/AI_WP_157.pdf?sequ

Horn, B. K. P. (1980). Derivation of invariant scene characteristics from images. *Proceedings of the May 19-22, 1980, National Computer Conference*, 371–376. ACM.

Hume, D. (1739). *A Treatise of Human Nature* (Reprint edition; D. F. Norton & M. J. Norton, Eds.). Clarendon Press.

Hume, D. (1777). *An Enquiry Concerning Human Understanding* (T. Beauchamp, Ed.). Clarendon Press.

Kant, I. (1996). *Critique of pure reason* (W. S. Pluhar & P. Kitcher, Eds.). Indianapolis, Ind.: Hackett Pub. Co. (Original work published 1787)

Kind, A. (2016). *The Routledge Handbook of Philosophy of Imagination*. Routledge.

Klüver, H. (1926). Mescal Visions and Eidetic Vision. *The American Journal of Psychology*, *37*(4), 502–515.

Klüver, H. (1928). *Mescal: the "divine" plant and its psychological effects*. London: Kegan Paul, Trench, Trubner & Co.

Klüver, H. (1966). *Mescal and Mechanisms of hallucinations*. Chicago: Univ. of Chicago Press.

Lamport, L., Shostak, R., & Pease, M. (1982). The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, *4*(3), 382–401.

Locke, J. (1689). *An Essay Concerning Human Understanding*. Oxford University Press.

Locke, J. (2012). *An essay concerning human understanding* (P. H. Nidditch, Ed.). Oxford: Clarendon. (Original work published 1690)

Macpherson, F., & Haddock, A. (2011). *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press Uk.

Macpherson, F., & Wilson, K. (n.d.). The Illusions Index. Retrieved March 26, 2019, from The Illusions Index website: https://www.illusionsindex.org/

Martin, M. G. F. (2002). The Transparency of Experience. *Mind and Language*, *17*(4), 376–425.

Martin, M. G. F. (2004). The Limits of Self-Awareness. *Philosophical Studies*, *120*(1-3), 37–89.

Martin, M. G. F. (2006). On Being Alienated. In *Perceptual Experience*. Oxford: Oxford University Press.

McDowell, J. (1983). Criteria, Defeasibility, and Knowledge. In *Proceedings of the British Academy, Volume 68: 1982* (pp. 455–479). Oxford University Press.

McDowell, J. (1986). Singular Thought and the Extent of "Inner Space." In J. McDowell & P. Pettit (Eds.), *Subject, Thought, and Context*. Clarendon Press.

McDowell, J. (1994). The Content of Perceptual Experience. *The Philosophical Quarterly*, *44*(175), 190–205.

Metzinger, T. (2003). *Being no one : the self-model theory of subjectivity*. Cambridge; London: The MIT Press.

Miskovic, V., Lynn, S. J., Bagg, J. O., Pouliot, J. J., Ríos, M., & Wackermann, J. (2019). Perceptual Phenomena in Destructured Sensory Fields: Probing The Brain's Intrinsic Functional Architectures. *Neuroscience and Biobehavioral Reviews*. https://doi.org/10.1016/j.neubiorev.2019.01.014

Nahm, F., & Pribram, K. (1998). Heinrich Klüver. *Biographical Memoirs. National Academy of Sciences* , *73*, 288 – .

Neumann, B. (2003). *A conceptual framework for high-level vision*. Retrieved from http://edoc.sub.uni-hamburg.de/informatik/volltexte/2009/55/

Pennartz, C. M. A. (2015). *The Brain's Representational Power: On Consciousness and the Integration of Modalities*. The MIT Press.

Perler, D. (2015). Perception in Medieval Philosophy. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception*. Oxford University Press.

Purves, D., & Lotto, R. B. (2003). *Why we see what we do : an empirical theory of vision*. Sunderlan, Ma: Sinauer.

Purves, D., Morgenstern, Y., & Wojtach, W. T. (2015). Perception and Reality: Why a Wholly Empirical Paradigm is Needed to Understand Vision. *Frontiers in Systems Neuroscience*, *9*, 156.

Rao, A. (2018, October 11). Our Experience of Reality Is a Bunch of Hallucinations We Collectively Agree On. Retrieved April 10, 2019, from Motherboard website: https://motherboard.vice.com/en_us/article/8xjbn3/consciousness-is-just-a-bunch-of-hallucinations-we-collectively-agree-on

Raz, G., & Seth, A. (2018, October 5). Anil Seth: How Does Your Brain Construct Your Conscious Reality? *NPR*. Retrieved from https://www.npr.org/2016/07/15/654730916/how-does-your-brain-construct-your-conscious-reality

Reid, T. (1785). *Essays on the intellectual powers of man*. Edinburgh.

Robinson, H. (1994). *Perception*. London: Routledge.

Robinson, H. (1996). Introduction. In G. Berkeley & H. Robinson (Eds.), *Principles of human knowledge ; and, Three dialogues* (pp. ix – xxxvii). Oxford; New York, N.Y.: Oxford University Press.

Russell, B. (1912). *The problems of philosophy*. London: Oxford Univ. Press.

Russell, B. (1948). *Human knowledge : its scope and limits*. New York: Simon and Schuster.

Searle, J. R. (2015). *Seeing things as they are: A theory of perception*. Oxford University Press.

Segall, M. H., Campbell, D. T., & Herskovits, M. J. (1963). Cultural differences in the perception

of geometric illusions. *Science*, *139*(3556), 769–771.

Seth, A. (2017). *Your brain hallucinates your conscious reality*. Retrieved from https://www.ted.com/talks/anil_seth_how_your_brain_hallucinates_your_conscious_reality?language=en

Siegel, S. (2008). The Epistemic Conception of Hallucination. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, Action and Knowledge* (pp. 205–224). Oxford University Press.

Simmons, A. (2015). Perception in Early Modern Philosophy. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception*. Oxford University Press.

Sloman, A. (1984). *New Horizons in Educational Computing* (M. Yazdani, Ed.). New York, NY, USA: Halsted Press.

Smith, A. D. (2002). *The Problem of Perception*. Harvard University Press.

Smythies, J. R. (1956). *Analysis of perception*. London: Routledge and Paul.

Snowdon, P. (1992). How to interpret "direct perception." In *The Contents of Experience: Essays on Perception* (pp. 48–78). Cambridge University Press.

Snowdon, P. F. (1980). Perception, Vision, and Causation. *Proceedings of the Aristotelian Society*, *81*(n/a), 175–192.

Strawson, G. (2003). What is the relation between an experience, the subject of the experience, and the content of the experience? *Philosophical Issues. A Supplement to Nous*, *13*(1), 279–315.

Strawson, P. F. (1970). Imagination & Perception. In L. Foster & J. W. Swanson (Eds.), *Experience & Theory* (pp. 31–54). University of Massachusetts Press.

Strawson, P. F. (1979). Perception and its Objects. In G. F. Macdonald (Ed.), *Perception and Identity: Essays Presented to A. J. Ayer with his Replies to them* (pp. 41–60). London:

Macmillan Education UK.

Swanson, L. (2016). The Predictive Processing Paradigm Has Roots in Kant. *Frontiers in Systems Neuroscience*, *10*. https://doi.org/10.3389/fnsys.2016.00079

Swanson, L. (2018). Unifying Theories of Psychedelic Drug Effects. *Frontiers in Pharmacology*, *9*, 172.

Williamson, T. (1990). *Identity and Discrimination*. Wiley-Blackwell.

Zucker, S. W. (1987). The emerging paradigm of computational vision. *Annual Review of Computer Science*, *2*(1), 69–89.

Appendix

The mechanisms of *S*-hallucination

In this Appendix I provide a technical account of how the objects of perception are produced by the mind with a set of mind/brain mechanisms which I think can make sense of veridical perception as well as non-veridical and partially-veridical illusions and hallucinations.

Stimuli

I use the word *stimuli* to refer to whatever it is that 'prompts' a perceptual experience. Importantly, my use of the term stimuli is intended to be totally generic and ontologically neutral. Stimuli can be internal or external; made of photons or vibration or chemicals or kinetic impact; or even electrical impulses coming from the supercomputers that power the Matrix. A stimulus is *anything* which can 'prompt' or 'drive' a perceptual experience, including the volition of an agent. So, for example: an actual ripe tomato sitting on a table, a hologram of a ripe tomato projected onto a table, a reflection of a ripe tomato against the surface of a mirror, or a brain implant capable of inducing a visual experience of a ripe tomato—all of these are *stimuli* which can prompt visual experiences.

Unknown causes and uncertainty

The cause of any stimuli that is not self-produced by the agent is unknown to the agent. Photons impinge on your retina and this activity is transduced into your nervous system. What caused those photons to impinge on your retina? From the vantage point of your

retina and by extension your nervous system, the causes are unknown, hidden, latent variables. This is the basis of the infamous "Brain in a Vat" thought experiment, where we are asked to imagine that we are not in the world as we experience ourselves to be but are instead simply brains suspended in a vat of nutrients with the appropriate stimuli impinging on our sensory inputs to create our experience of being in the world. "Is it possible to imagine such a thing? I think that it clearly is, for a very simple reason: we are in fact brains in vats. The skull is a cranium vat made largely of calcium, but it is roughly spherical and it houses a brain" (Searle, 2015, p. 77). Similarly, this fact is what Hohwy (2013, see also 2016) terms "the skull-bound brain" and what Clark (Clark, 2013, p. 183) characterizes as the "view from inside the black box"—a problem well-known to even to modern-period thinkers from Hume to Kant to Helmholtz (Swanson, 2016). In summary, all stimuli have *unknown causes*.[26]

The Byzantine generals problem

A nice way of characterizing the situation of being a perceiving mind comes from the computer science literature and is known as the "Byzantine Generals Problem" (Lamport, Shostak, & Pease, 1982). This is a way of describing the challenge faced by certain types of systems, typically distributed computer systems, but here I apply it to perceiving minds (brains). "Reliable computer systems must handle malfunctioning components that give conflicting information to different parts of the system" (Lamport et al., 1982, p. 382). Lamport et al. use a metaphor to describe the situation:

> We imagine that several divisions of the Byzantine army are camped outside an enemy city, each division commanded by its own general. The generals can communicate with one another only by messenger. After observing the enemy, they must decide upon a common plan of action. However, some of the generals may be traitors, trying to prevent

[26] Yes, this echoes the Hume-inspired Kantian realization that things-in-themselves are not known to the agent.

> the loyal generals from reaching agreement. The generals must have an algorithm to guarantee that
>
> A. All loyal generals decide upon the same plan of action.
>
> The loyal generals will all do what the algorithm says they should, but the traitors may do anything they wish. The algorithm must guarantee condition A regardless of what the traitors do.
>
> The loyal generals should not only reach agreement, but should agree upon a reasonable plan. We therefore also want to insure that
>
> B. A small number of traitors cannot cause the loyal generals to adopt a bad plan (Lamport et al., 1982, p. 383).

The authors point out that the problem *seems* "deceptively simple" yet is in fact tremendously challenging. Here, I suggest that this is the central problem that brains evolved to solve[27]. It is important to understand my perspective on this as it has motivated my *S*-hallucination theory. Imagine that you are in the BWCA being chased by a murderous bear (or a murderous man, which is the more likely scenario). As you sprint down the rugged, obstacle-ridden trail, your nervous system—a distributed system with many disparate components—needs to solve something akin to the Byzantine generals problem in order to navigate and evade the predator. It needs to act as a unified system while the individual components are sending and receiving conflicting and even erroneous messages. The available sensory information is limited and imperfect: you see and hear the predator behind you, you see and hear and feel the obstacles underfoot and just ahead of you, and all of this data must be weighed and integrated to inform your every move. Imagine now that you (think) you have evaded your predator, but you are not sure. You think you hear the raspy growl of your predator to your left, but you (think you) see a silhouette of him to your right. Your visual system and your auditory system are sending conflicting messages. Which one should guide the system? Which one is 'veridical'? You don't know for sure, so you have to make your *best guess*.

[27] So far as I am aware, I am the first to explicitly apply the Byzantine Generals Problem to describe the dynamics of perceptual processes in biological systems.

Causal inference

Hume pointed out that, since the causes of stimuli are unknown to the perceiver, perception must employ some kind of inferential process in order to deliver the meaningful connections presented in a perceptual scene (Hume, 1739, bk. bk. I, part III, section vi). This process has come to be known as *causal inference.* Given that stimuli have unknown causes and that the brain must solve a Byzantine generals problem, it makes sense to characterize the structure of your subjective perceptual field as the product of your mind's 'best guess' as to the causes of its driving stimuli. Luckily, it is an *informed* guess: it is based on statistical regularities.

Statistical regularities

How can causal inference transform stimuli with unknown causes into a complete subjective perceptual field rich with structure and meaning? While the *causes* of stimuli are unknown, many stimuli nonetheless occur with structured regularity. Stimuli are usually not random noise. The statistics of these regularities are available—and actively exploited—to bootstrap causal inference into generating the objects, entities, and events presented in the scenery of your subjective perceptual field. Regularity entails sampling over time, and for the samples to evidence regular patterns, they must be stored in some manner so as to provide ongoing availability to the agent.

Beliefs and prior knowledge

Prior knowledge is the result of the agent tracking and storing statistical regularities from past experience. Prior knowledge ('priors' for short) can also be described as internal

networks of (mostly unconscious) *beliefs* about the workings and states of affairs of the world; the causes of stimuli. Causal inference recruits beliefs to explain the unknown causes of impinging stimuli.[28]

Generative schemata

> Our brains build models of the world and continuously modify these models on the basis of the signals that reach our senses. So, what we actually perceive are our brain's models of the world. They are not the world itself, but, for us, they are as good as. You could say that our perceptions are fantasies that coincide with reality (Frith, 2007, p. 135).

A term borrowed from Kant, a *schemata* refers to the entire apparatus internal to an agent which generates perceptual scenes for that agent using causal inference powered by a combination of beliefs and stimuli. The machine learning technique known as *generative models* is roughly akin to the Kantian idea of schemata (Swanson, 2016). A generative model "aims to capture the statistical structure of some set of observed inputs by tracking (one might say, by *schematically* recapitulating) the causal matrix responsible for that very structure" (Clark, 2013, p. 182 emphasis mine).

So your brain is a model of the regularities in stimuli, and your subjective perceptual experience is generated via that model. Importantly, *this* is the origin of the 'objects' in your perceptual field. The objects of perception are not 'given' in the stimuli; rather, they emerge as products of your mind's guess at what is causing the stimuli.

---

[28] This section describes in simple terms what is a more complicated set of technical concepts. In Bayesian statistics, beliefs are probability distributions over unknown states or attributes. Beliefs have associated *precision*, which is the inverse variance or the most likely value (mean) over which the probability distribution is concentrated, also known as the *expectation*. Precision quantifies the confidence or certainty associated with a belief. *Prior* beliefs exist before seeing the data, which are updated to *posterior* beliefs after seeing the data—a prior belief is transformed into a posterior belief when it is combined with sensory evidence or the *likelihood* of the data. When we treat expectation and precision themselves as unknown quantities, we get beliefs *about beliefs*, often called *empirical priors*, such as when we hold an expectation about a precision associated with a belief. When Bayesian inference is deployed in a hierarchical fashion, beliefs about beliefs are inevitable as they provide constraints on lower-level beliefs. The process of forming a belief about the precision of a belief is sometimes called *precision-weighting*.

Working together to generate a perceptual scene

With the above terms introduced[29], we can run through the full *S*-hallucination account of perceptual experience. Consider that your subjective perceptual field is driven by stimuli with unknown causes, yet still presents a complete scene made up of phenomenal wholes which are composed from their smaller composite parts to form the overall Gestalt of the scene. Looking out my window at dusk, for example, I don't experience formless blobs of red and green and black—instead, "I see the red light of the setting sun filtering through the black and thickly clustered branches of the elms; I see the dappled deer grazing in groups on the vivid green grass …" (P. F. Strawson, 1979, p. 43).[30] Perceptual scenes are rich with content, semantic meaning, intentionality.

How does such a scene arise from tiny bits of impinging sensory stimuli coming from unknown causes? On the view I am articulating here, every object, entity, background, foreground, and event in your subjective perceptual field is generated by internal mechanisms which infer the causes of the stimuli. How does your mind do this? It tracks regularities in its stimuli. It imposes *meaning* or *aboutness*—intentionality—on these regularities and stores this meaning as internal (generative) models of the causes of the stimuli. Your mind uses the generative models to produce a subjective perceptual field containing *objects*: groupings, categories, generalities. Due to the fact that the causes are hidden and uncertain, the process is inherently probabilistic and inferential—a best guess.

But, you say, it doesn't *feel like* a guess. It seems like the world, the objects, entities, and events in the scene, are just there, plain as day, freely available for the

[29] The reader with a background in Bayesian neuroscience or Predictive Processing theories will recognize these core concepts. Here I aim to provide the most generic high-level conceptual framework possible for *philosophical* purposes, free of any specific neuroscientific implementation.

[30] True story: as I wrote this paragraph, I looked out my window and actually saw deer grazing at dusk—a veridical perception.

taking, full of certainty. While perception may subjectively feel certain and clear, this does not refute the claim that it is a best guess.