The Sense-Data Language and External World Skepticism

Jared Warren

(penultimate draft – please cite and quote from the final version, to appear in the *Oxford Studies in Philosophy of Mind* 4)

Abstract

We face reality presented with the data of conscious experience and nothing else. The project of early modern philosophy was to build a complete theory of the world from this starting point, with no cheating. Crucial to this starting point is the data of conscious sensory experience – *sense data*. Attempts to avoid this project often argue that the very idea of sense data is confused. But the sense-data way of talking, the sense-data language, can be freed from every blemish using ideas from contemporary metaontology. We can adopt a sense-data framework that vindicates the traditional claims of sense-data theories and leads to plausible theories of perception and color. We can, we should, and in a sense, *we must*. Yet when we do, we face the traditional problem of external world skepticism, head-on. The real challenge of skepticism is forced upon us; it cannot honestly be avoided using externalist tricks, burden of proof shifting, or other razzle dazzle. But the challenge can be met: we are rational to posit the external world as the best explanation of the many synchronic and diachronic patterns over our sense-data. This paper argues for all of these points and ends with a plea for *analytic empiricism* – a traditional sense-data version of empiricism that uses all of the tools of analytic philosophy while avoiding the more questionable doctrines of the Vienna Circle (phenomenalism, verificationism). Analytic empiricism is our last best hope for completing the great philosophical project of early modern philosophy.

I. The Early Modern Project

Philosophy starts with the data of consciousness and builds outwards. At least, that's what the early modern philosophers did, starting with Descartes. This starting point allows us only what is directly accessible to consciousness. The real problem of external world skepticism is a problem raised from this starting point. Because of this, modern skepticism goes beyond the skepticism of the ancients.[1] This goes for every school of ancient skepticism — the Pyrrhonians, the Academic Skeptics, and the Cārvāka.

The early modern project, inaugurated by Descartes, was to build a complete theory of the world outward from this starting point. In many ways, the story of contemporary philosophy is the story of a flight from Descartes. It is only slightly misleading to say that both analytic and continental philosophy were created by this flight. The founders of these traditions are often credited with reorienting philosophy. Reorienting it, that is, away from the first-person starting point.

Depending on what you're interested in, reorientation might be advisable, but is it *required*? Nearly everyone seems to think so. There is a very — *very*! — widespread feeling that the early modern project is broken, cracked to its core. I disagree. The first-person starting point remains viable. What is more, the project is, in a real sense, forced upon us. What is even more, there is yet hope that we can complete the project without cheating. All of this I will argue for here. To say that I have my work cut out for me is an understatement.

II. From Ideas to Sense-Data

The central crack in the early modern project is supposed to be in its characterization of the *data* of consciousness. For Descartes, Locke, Berkeley, and others this data was said to consist of "ideas". Hume departed from this in terminology but not in doctrine by using "ideas" only for the less vivid mental particulars and "impressions" for the others.

This "idea" terminology was extended to cover all mental particulars. Not just mental images, not just data from other senses, but even emotional states. This already stretched the terminology too far, but the early moderns stretched it further still. They also tried to

[1] See Burnyeat 1982 and Williams 2010.

understand *beliefs* and *concepts* and *meanings* using ideas.[2] This is a hopeless approach to cognition and we should have no part of it. To free ourselves from it, both terminology and theory need updating.

The early moderns themselves might have gone along with this — we already find Berkeley writing:

> As for our senses, by them we have the knowledge only of our sensations, ideas, or those things that are immediately perceived by sense, call them what you will.[3]

Call them what you will; I will call them sense-data (plurally) or sense-datum (singularly). This terminology was popularized by Moore and Russell in the early twentieth century.[4] Subsequently the notion of sense-data, if not always the terminology of "sense-data", was widely accepted by British philosophers.[5]

Sense-data are supposed to be what we are directly *aware of* or *conscious of* during perception. The mere existence of hallucinations and illusions shows that conscious, perceptual experiences do not require the existence of external objects.[6] As the terminology suggests, *sense-data* covers less ground than *ideas*. It includes only the data of sensory experience. Sense-data are private objects accessible to consciousness and usually associated with a single sense.[7] They are not *inherently* intentional or representational, but are instead phenomenal or qualitative.[8] With sense-data there is no significant gap between

---

[2] For discussion, see section 17 of Bennett 1966.
[3] From Section 18, Part I of *A Treatise Concerning the Principles of Human Knowledge* — see page 89 of Berkeley 2008.
[4] Perhaps the first published use of the term is in Moore 1909-10, it was popularized by Russell 1912. See also Moore 1913-4 and Russell 1914*a*,*b*.
[5] For prominent examples from the period, see Broad 1923, Price 1932, and Ayer 1940.
[6] Arguments of this kind are offered in, for example, Broad 1923 and Price 1932.
[7] But see the final chapter of Ayer 1968 for a useful discussion of whether privacy should be stipulated, up-front.
[8] As I will discuss in section 6, these lines start to blur at a few points.

appearance and reality. Some proponents also stipulate up-front that sense-data are non-physical, but I think their metaphysical nature should initially be left open.[9]

In some ways, the "sense-data" terminology is also less than ideal. It might suggest *raw data* from the senses, but that isn't right. Sense-data are instead the sensory data *accessible to first-person conscious experience*. For this reason, Ayer later switched to "sense-qualia" instead of "sense-data".[10] The reasoning behind this terminological change is sound, but, unfortunately, the "qualia" terminology has even more baggage. Russell also later adopted a different terminology, in his case moving to "percepts". He also characterized this move as a *rejection* of sense-data, but this is misleading.[11] I agree with Ayer's assessment:

> …[Russell's] rejection of sense-data turns out only to be a rejection of the subject-object analysis of sensation.[12]

So despite their respective moves away from "sense-data" terminology, both Ayer and Russell remained sense-data theorists until the end.[13] Despite some minor misgivings, I'm going to stick with the "sense-data" terminology that they (and Moore) popularized.

Sense-data enjoyed a brief heyday from the early to mid-twentieth century. But by mid-century the very idea of sense-data was under heavy fire. The most damaging assaults were launched by the towering ordinary language philosophers, Ryle and Austin.[14] A consensus *against* sense-data started to build, a consensus that the battle against sense-data had been well-fought and fairly won. By the late 1970s Charles Taylor could open a critical article on sense-data by saying:

---

[9] Robinson (1994) stipulates that sense-data are non-physical.
[10] See Ayer 1973, for example.
[11] See the discussion in Russell 1959.
[12] Ayer 1969, page 174.
[13] Many other sense-data theorists also used different terminology – for example, Broad 1923, Price 1932, and Mackie 1976.
[14] See Ryle 1954 and Austin 1962.

> I want in this paper to return to a subject which has been out of fashion now for about two decades, the subject of sense data.[15]

So over forty years ago sense-data had already been "out of fashion" for *twenty years*. Today the subject is covered with more than *sixty years* of dust and cobwebs.

Even under this debris, the sense-data theory is not totally dead, but it is now extremely unpopular. In the 2020 philpapers.org survey not even 1% of surveyed philosophers accepted sense-data, and only 4% "leaned toward" accepting sense-data. In this environment, sense-data cannot simply be assumed the way the early moderns assumed the existence of ideas. For them, the idea framework was the starting point. They didn't think it required a significant defense. But after criticisms by Reid and Kant, ideas could never again be assumed without comment. Likewise now, for sense-data. Accordingly, contemporary defenders treat sense-data theories as *substantive* positions.[16] They think we can only adopt sense-data talk at the *end* of the trail, not at the *start*. With many intricate arguments and rebuttals and constructions required to get us to that point.

I think that this is a mistake. The move to sense-data is the move to a sense-data *language*.[17] It is a way of approaching experience, not a substantive theory about the nature of experience. The sense-data language is extremely useful. It is useful for talking about experience, for posing philosophical questions, and for offering theories that demystify the apparently mystical. Only within this sense-data language can substantive sense-data-based theories of perception be given. Yet the existence of sense-data does not hinge on the success of such theories. As Descartes and Locke and Berkeley and Hume all assumed, we can talk about sensory experience on the cheap, for the sense-data language is their idea language shorn of overreach and muddles. To defend these claims, I'll draw from contemporary metaontology. As a lead-in to that, let me introduce sense-data talk in more detail.

---

[15] Taylor 1979, page 99.
[16] In particular, I'm thinking of the major defenses in Jackson 1977 and Robinson 1994.
[17] Grammatical purists might prefer calling it "the sense-datum language", but to my twenty-first century English ears, "the sense-data language" now sounds more natural.

III. A Framework for Sense-Data

I see a red apple on the counter in front of me. Given this, a red apple exists. Red apples aren't in my mind, so it follows that *something* — something outside of my mind — exists. This reasoning won't satisfy the skeptic, because sometimes we *seem* to see things without *actually* seeing them

We can't always tell the difference, from the inside, between seeing an *F* and seeming to see an *F*.[18] This cannot be reasonably denied.[19] Let's carefully go over how sense-data talk is introduced on this basis.[20] Start with a standard natural language report about visual experience:

(1) I see a red apple on the counter in front of me

This entails that there is a red apple on the counter in front of me. Yet, as just noted and as stressed by historical sense-data theorists, we sometimes think we see something in the world when we actually do not. Much more secure than (1) is the related claim about how things seem to me:

(2) It seems to me that I see a red apple on the counter in front of me

On one reading, (1) does not entail (2) — I might see a red apple that looks brown because of the lighting. And obviously (2) does not entail (1) either — I might be hallucinating a red apple.

[18] A sentence like "I seem to see an apple" is here equivalent to: I am currently having a visual experience that is as if I'm seeing an apple.

[19] Cases of change blindness show that we sometimes judge that scenes look the same without them actually looking the same, yet on reflection, the scenes can be distinguished. This isn't relevant to the claim here.

[20] A similar approach to the introduction of sense-data talk is taken in Ayer 1956.

The distinction between (1) and (2) is uncontroversial. Yet it provides the foundation on which the sense-data language is built. To start talking about sense-data explicitly we add a transition from (2) to:

(3) I see a seeming-red-apple-on-the-counter-in-front-of-me

Where (3) is used in an ontologically committing way. Given a standard, Quinean account of ontological commitment, this means that from (3) we infer:[21]

(4) There is something that I see

Of course, the witness to (4) can't be an ordinary object since, as already noted, (2) can be true even when no external object is being seen. (2) is supposed to entail (3) which is supposed to entail (4). Given these entailments, (4) can be true when no external object is seen.

The ordinary language critics of sense-data seized on this point. They argued that by talking like this we sin against the ordinary usage of "see".[22] My example concerns vision, but similar points also go for the other canonical senses. As a point of descriptive linguistics there might be something to this objection. Speaking with the vulgar, I *see* a red apple. Speaking with the learned, I *see* a seeming-red-apple. The red apple is not the seeming-red-apple. Therefore, speaking with everyone, there are at least two distinct things that I see, in this very moment, in this very act of perceiving. That sounds extremely strange.[23]

[21] From Quine 1948.
[22] The criticism is most associated with Austin 1962, but a similar point had already been made about "perceives" in Reid's 1785 *Essays on the Intellectual Powers* — see Reid 1983. But Reid was more subtle (or more fair-minded) than Austin, since he also admitted that we are "conscious of" ideas (sense-data). For a contemporary endorsement of Reid's criticisms, see Maddy 2017.
[23] Mates (1981) makes a similar point.

We could live with the strangeness if we had to, but it's more comfortable to introduce special terminology for the way in which we "see" a sense-datum.[24] We could unimaginatively call this "see*", but it's slightly more artful to instead replace (3) and (4) with the following principles:[25]

> (3*) I am directly visually aware of a seeming-red-apple-on-the-counter-in-front-of-me
>
> (4*) There is something that I am directly visually aware of

The sense of "direct visual awareness" here is purely *phenomenal*. It applies only to sense-data, so there is no double counting. We *see* ordinary objects but are only *directly aware* of sense-data.

Versions of this terminological move have a long history. I already quoted Berkeley introducing the terminology of "immediately perceiving" to the same purpose. Others have said that we "directly" see only sense-data.[26] Still others that we are "acquainted" or "directly acquainted" only with sense-data.[27] Whatever we call it, this notion of *direct acquaintance* is crucial. It will be spelled out further below.

This terminological change seems cosmetic, but even by itself it goes a long way toward side-stepping ordinary language complaints about sense-data.[28] With this change made sense-data theorists no longer need to say — implausibly — that we do not ever *see* trees, travelers, and trebuchets. We do see ordinary objects, often and happily. The sense-data theorist instead says that while we see trees, we are never directly visually aware of them.

---

[24] This also side-steps objections from the premise that *in English* "see" is not always ontologically committing.
[25] Paul 1936 replaces "looks" with "looks[1]" for roughly the same purpose.
[26] Jackson 1977 says this.
[27] Fumerton 1985 and Russell 1912 say something like this.
[28] It also side-steps worries about the causal relationship between *us* and *our sense-data*. We can think of direct acquaintance as constituted, in part, by sense-data without requiring a causal connection *between* sense-data and an independent self. This is related to Ayer's claim quoted above, against Russell, that rejecting a subject-object analysis of sensation is not tantamount to rejecting sense-data.

We are never directly acquainted with them. This is the only conclusion that is established by the infamous arguments from illusion and hallucination, not any stronger conclusions about "seeing" in the ordinary sense.[29] We can accept that sense-data is all we are ever "acquainted with" or "directly visually aware of" in this *special* sense, but not that we only ever "see" sense-data.[30] If you also want to say that there is some way we "directly" see trees, feel free. Just take care to distinguish *that* sense of direct perception from *this one*.

When we reason in this way to (3*) and (4*) we have already introduced a sense-data extension of natural language. The extension involves "direct awareness" relations for each of our senses and a new family of terms, ⌜ seeming-*F* ⌝; for any natural language noun phrase *F*. To this we also add an explicit sortal predicate, "sense-datum", governed by the following principles:

If something is a seeming-*F*, then it is a sense-datum

And:

Someone is directly aware of α if and only if α is a sense-datum

This biconditional principle can be understood tenselessly. In adopting its right-to-left direction, I am ruling out — as *existing* — objects of the same metaphysical status as sense-data that are unsensed over all of history, what Russell called "sensibilia".[31] I am also implicitly ruling out the identification of sense-data with aspects of ordinary objects, such

[29] Searle (2015) calls the stronger version of the argument from illusion "the bad argument" and he blames it for many horrors that philosophy has delivered over the ages. Robinson (1994) argues that the role of the argument from illusion and related arguments in the history of philosophy has been overblown. Hirsch (2002) notes that, if you argue yourself into accepting "we never see trees", you will have unwittingly altered English and ended up speaking an alternative language.

[30] For some relevant empirical information concerning whether English speakers would *say* that we "see" or "perceive" hallucinations, see D'Ambrosio 2022.

[31] See Russell 1914*a*. Note however, that we can still talk about merely possible sensory experiences, see sections 7 and 11 for relevant discussion.

as their surfaces.[32] Neither of these exclusions is based in *metaphysics*. I am merely specifying what *I* mean by "sense-data", while keeping my usage firmly within the mainstream among historical sense-data theorists.

In introducing these adjuncts to English I have characterized what Carnap called a "linguistic framework" for talking about sense-data.[33] A few of the second-wave sense-data theorists explicitly recognized that the sense-data-framework had to be distinguished *linguistically* from the ordinary-object-framework.[34] This contrasts with the procedure of the most prominent recent sense-data theorists, who think the introduction of sense-data is both philosophically and empirically substantive.[35]

To fully flesh out and defend this linguistic approach and its philosophical relevance, we need to reflect carefully on the metaontology of the sense-data framework.

IV. The Metaontology of Sense-Data Talk

Suppose we start with English and then add sense-data vocabulary governed by the inferences from (2) to (3*) and from (3*) to (4*) . We thereby start speaking a new extended language, SD-English — English augmented with the ability to talk about sense-data. I claim that SD-English is a coherent language.

I mean this in a strong sense. Everyone agrees that we could start speaking SD-English, in that we could reason using the inferential transitions from (2) to (3*) and the like. My claim goes beyond this near platitude. I additionally claim that *if* we started to speak SD-English in this benign sense, *then* the inferential transitions from (2) to (3*) and from (3*) to (4*) would be *valid in our new language*. If it is true that I seem to see a red apple, this means the following sentence will also be true:

---

[32] See Ayer 1945 for defense of similar stipulations against Moore 1918-9.
[33] In Carnap 1950.
[34] Paul 1936 and Ayer 1940, 1945 are particularly clear examples of this recognition. Also, see chapter 3 of Mates 1981 and the brief discussion on pages 60-1 of Hirsch 2002.
[35] As mentioned above, here I have in mind the major defenses in Jackson 1977 and Robinson 1994.

There is a seeming-red-apple

And since a seeming-red-apple is a sense-datum, this means it will also be true that sense-data exist. True *in SD-English*, that is, not in English. These aren't sentences of (pre-philosophical) English, so the question of their truth in English does not even arise.[36]

Your spider senses are probably tingling right now. It *sounds* like I'm claiming that merely by *talking* in a certain way, we *create* objects. But that's silly. We can only create sound-waves and word-tokens and the like by talking. Whatever sense-data are, they aren't created just by speaking a language. So if I were making this claim, you would be right to worry. But I am not. Clarity on this point has been hard-won, only worked out carefully over the last twenty-five years of developments in metaontology. Let's go over it in detail.

In full generality, the step from (2) to (3*) looks like this:

I seem to see an *F*

Therefore:

I am directly visually aware of a seeming-*F*

We can allow this to apply even when *F* isn't obviously sensible — "I seem to see some kosher food" — but doing so will introduce a more than moderate amounts of indeterminacy (see the discussion in sections 6-8).[37] This transition is an example of what Schiffer called a "something-from-nothing transformation".[38] Such transitions move from a sentence that seems ontologically neutral with respect to some objects to one that is ontologically committed to those objects. But despite the ontological commitment, the transformation

[36] English talks about some particular types of sense-data — "I saw a red afterimage" — and may even talk about sense-data generally with terms like "appearances". I am not denying this, of course, and it will be relevant in sections 7 and 11 below.
[37] Thanks to Eli Hirsch here.
[38] In Schiffer 1996, 2003.

itself also *seems* completely unobjectionable. It even seems to — in a coarse-grained sense — preserve meaning.

Something-from-nothing transformations are all over the place. Schiffer focused on the following three kinds:

(*PROPERTIES*) the ball is red, therefore: the ball has the property of redness

(*NUMBERS*) there are four moons, therefore: the number of moons is four

(*PROPOSITIONS*) Lassie is a dog, therefore: the proposition that Lassie is a dog is true

The idea that something-from-nothing transformations or the like can settle ontological debates has been dubbed "easy ontology", by Sider.[39] The most fully worked-out version of easy ontology is probably Neo-Fregeanism in the philosophy of mathematics.[40]

The central idea behind easy ontology is that the crucial something-from-nothing transformations are analytic or conceptual truths. This invites controversy. Not only because the very existence of analytic and conceptual truths has been prominently denied.[41] But also because the idea that *existence claims* could be analytic, *even conditionally*, seems especially hard to fathom.[42]

Making sense of it requires a metaontology — a meta-theory for ontology — showing how existence claims can be cheap. Sider, Hirsch, and I have all argued that easy ontologists must be implicitly assuming *quantifier variance*.[43] The name, "quantifier variance", comes

[39] In Sider 2011.
[40] See Dummett 1956, Wright 1983, Hale 1987, and Hale and Wright 2001. Easy ontology is also endorsed in Thomasson 2009 and Fine 2005.
[41] See Quine 1951 and Williamson 2006 for the denials, and Grice and Strawson 1956 and Warren 2021*b* for rebuttals.
[42] A nice statement of this point of view is found at the start of Field 1984. See also the discussion in chapter 8 of Warren 2020.
[43] See the discussions in chapter 9 of Sider 2011, the introduction to Hirsch 2011, and chapter 9 Warren 2020. There has been opposition to this — see Hale 2007, Hawley 2007, and Eklund 2006 — but taking up this debate here would be a distraction.

from Hirsch, who has done more than anyone else to develop, defend, and apply the view.[44] But the doctrine is older than the name. Quantifier variance was previously advocated by Putnam and perhaps also by Carnap and others.[45]

According to quantifier variance there are different possible ontological languages. These are languages in which the respective unrestricted existential quantifiers, as individuated by syntactic inferential role, *mean different things*. The classic illustrating example involves mereology. In a world with two mereological atoms, *a* and *b*, located at different places in an otherwise empty space, how many *objects* are there? How many *things* exist in this world? In English, we say that only two things — *a* and *b* themselves — exist in that world. Yet it is easy to imagine a language just like English save for including the principles of classical mereology. Among these principles is a universal composition axiom holding that for any two distinct objects there is another object with each original object as a proper part. In this alternative language it is then correct to say that there are three things in the imagined world — *a*, *b*, and also their mereological fusion, *a*+*b*.

This is an apparent disagreement about *how many* objects exist. Therefore, it cannot come down to a disagreement about how predicates are distributed over given objects.[46] We have two distinct and seemingly incompatible answers to a straightforward ontological question about a simple toy case. Which answer is correct? Metaphysical realists assimilate this ontological disagreement to standard factual disagreements. There is supposedly a matter of objective fact here that the two parties in the dispute differ over. Many of us find this assimilation strange. What possible matter of fact is hidden here? In our toy case all of the facts are in. I completely described everything about the possible world in question. Nothing about the case was hidden. Nothing about the case could be hidden. Rather than a factual difference, it is more plausible to see this dispute as owing to a difference of *language*. The mereologist speaks a language that differs from ours. In particular, their language differs from ours over the meanings of "there is", "exists", "objects", "things", and related expressions.

---

[44] See especially the essays in Hirsch 2011.
[45] See Putnam 1987, 2004, and Carnap 1950.
[46] It also can't come down to only the identity predicate; see the appendix to Warren 2022*a*.

The background argument for this appeals to charity in metasemantics. Following certain rules and using language in certain ways, with a clear head, guarantees that the meanings of your sentences vindicate said uses. This is the basic claim of constitutive metasemantic charity.[47] Essentially the same point can also be made by appeal to the slogan that "meaning is use", or in terms of conceptual or inferential role semantics (inferentialism).[48] On all of these approaches, because of how the mereologist is disposed to use expressions like "there are" and "exists", he speaks *truly* when he describes the *a*-*b* world as a world where there are "exactly three objects". And because of how we are disposed to use language, we speak truly when we describe the *a*-*b* world as a world where there are "exactly two objects".

The two languages are quite similar, but they differ over the precise use of quantifier expressions. Even still, the structural features of the use of these quantifier expressions are common to both languages. In particular, in both mereological and non-mereological English, the expressions "exists" or "there are" play the exact same structural inferential role. That is: in both languages, these expressions play the inferential role of the unrestricted existential quantifier familiar from the elementary logic classroom. According to the standard, Quinean account of existence and ontological commitment, both languages have a concept of existence that is expressed by their unrestricted existential quantifier. Using their own existential quantifier, they can make existence claims. Yet because of subtle differences in use, the two languages have *distinct* concepts of existence. The meanings of the quantifiers and related expressions *vary* across the two languages.

This is the crucial point. In neither language are any objects being "created" by mere talk. The linguistic activities of language users instead only determine what is meant, which concepts are expressed. The best way to describe the situation is by saying that the languages both have concepts of existence, but that the two concepts of existence differ. Despite the differences, they both still count as concepts of *existence*, or as existence-like concepts, because of their relevant inferential similarity. They are similar, they are both unrestricted quantifiers, yet they are distinct. Sentences containing these expressions can

[47] Charity was introduced in Wilson 1959 and was influentially advocated by Quine (1960*a*), Davidson (1973), and Lewis (1974), among others.
[48] The slogan is associated with Wittgenstein (1953).

differ in their truth values across the two languages. This example shows that our familiar English concept of the existence of an object is not built into all possible languages.

This is roughly the claim of modest quantifier variance. Strong quantifier variance adds to the pluralism of modest variance the claim that some pairs of distinct quantifier languages are expressively equivalent and therefore of equal metaphysical merit.[49] So for the on-record quantifier variantists – Putnam, Hirsch, myself – metaphysical merit is understood in terms of expressive power.[50] Expressive equivalence is plausible in the mereological case — any world you can describe in English can just as well be described in mereological English, and vice-versa. Hirsch has argued that many — though not all — of the ontological debates of philosophers satisfy the conditions of strong quantifier variance. That is: the ontological languages associated with the opposing sides in the debate are expressively equivalent and therefore (according to variantists) of equal metaphysical merit.

Ontologists reject this account of metaphysical merit. The best alternative theory, from Sider, accepts modest variance but understands metaphysical merit using a distinctive metaphysical primitive – *metaphysical structure*.[51] This way of rejecting strong variance allows the coherence of the sense-data language. But it enables critics to deny the existence of sense-data in the metaphysically special sense of "existence". For my part, I don't think that metaphysical structure is coherent; and even if it were coherent, I don't think it could have a reasonable epistemology.[52] For both of those reasons, I reject the concept of metaphysical structure and all related concepts.

Quantifier variance has, of course, faced many related and further objections from metaphysicians.[53] I think that it can be defended from these objections. In fact, I think it *has been* defended from these objections.[54] I won't recapitulate all of the back-and-forth here. My goal right now is to apply quantifier variance, not defend it. My point is that the application of quantifier variance vindicates the sense-data language.

---

[49] For both modest and strong variance, see Hirsch and Warren 2019*a*.
[50] See Putnam 1983, Hirsch 2011, and Warren 2015.
[51] See Sider 2009, 2011.
[52] See Warren 2016*a*,*b*.
[53] For examples, see Dorr 2005, 2014, Eklund 2009, and Hawthorne 2006, 2009.
[54] See Hirsch 2011, Hirsch and Warren 2019*a*,*b*, and Warren 2015, 2022.

So turning back to sense-data, note that the same charity argument that applied to the imagined mereological language also applies to SD-English. Speakers of SD-English clear-headedly accept the inferential transitions from (2) to (3*) and from (3*) to (4*) . In so doing they follow general rules of inference for the newly introduced sense-data terminology. Given that they clearheadedly use language in this fashion, charity requires us to interpret them as speaking a language in which these inferential transitions are valid. So, since it is true that we seem to see things, true as a matter of empirical fact, then it also must be true that sense-data exist — true in SD-English, that is.

Well-known issues of social externalism might block English speakers from transitioning to SD-English, but those issues can be side-stepped.[55] I hereby stipulate that our imagined SD-speakers do not explicitly or implicitly defer to English speakers in their relevant uses of language. If SD-English extends English, then its expressions, "there is", "exists", "object", and "thing" differ slightly in meaning from their English counterparts while retaining the same structural-inferential roles. SD-English speakers have a slightly different concept of *existence* than standard English speakers do.[56] English and SD-English are an example of modest quantifier variance. Below (section 7) I will argue that they aren't an example of strong quantifier variance, since SD-English has more expressive power than English (at least if English is stripped of its existing sense-data talk).

The quantifier variance metaontology vindicates SD-English's something-from-nothing transformations and — given the facts of sensory experience — the truth of the sentence "there are sense-data". A few earlier sense-data theorists approached matters from a similar perspective, but without the supporting metaontological details.[57] And the canonical early moderns took the "idea" language in a similar spirit, seeing it as a background framework

[55] See Burge 1979.

[56] I haven't said anything about the other famous Quinean metaontological slogan — "no entity without identity" from Quine 1957-8. But a criterion of identity for sense-data poses no special problems, for we can simply appeal to Leibniz's law. It is true that different sense-data theorists may differ over the properties of sense-data, but most of these are best seen as questions of *decision*, not *discovery*. Paul (1936) nicely advocated for a similar point of view. See the relevant discussion in section 5 and 7 below.

[57] In particular, again see Ayer 1940, Paul 1936, and Mates 1981.

that did not require substantive defense. Quantifier variance frees these approaches from every blemish. It is not that sense-data are created by language. Rather, the sense-data language talks about experience in an ontologically committing way, using its distinct concept of existence. If you speak the sense-data language, the standard claims about sense-data and their existence are *true*, unproblematically so, without sense-data being creations out of language.

From this perspective, objections to the existence of sense-data begin to seem a bit silly. They seem as silly as objections to the existence of tables seem to many of us. And the complex alternative schemes for talking about the data of sensory experience in a non-ontologically committing way start to seem, not silly, but unmotivated.[58] Both metaphysical and epistemological questions about sense-data are illuminated by this language-based way of looking at things. Let me demonstrate this.

[58] I have in mind the adverbial theory, advocated in Chisholm 1957. Adverbial theories analyze experience claims like, "I see an orange circle", using adverbs that modify the perceptual verb, for instance, "I see orange-circle-ly". This approach is sometimes endorsed over sense-data on grounds of parsimony, but the gain in ontological parsimony here is paid for in the coin of ideological profligacy — compare: nihilists who reject chairs but talk of simples arranged chairwise, following van Inwagen 1990. There is a famous problem for the adverbial theory, the so-called "many properties" problem — Jackson 1975, Casullo 1987, and Dinges 2015. Here is the problem: suppose I see an orange patch and circular shape. To simply say, as before, that I see orange-circle-ly is wrong. There is a difference between seeing something orange and seeing something circular at same time when they are distinct, versus seeing something that is both orange and circular. Object talk is our standard way of making exactly these distinctions, so the sense-data alternative faces no analogous problem. Adverbialists have attempted to solve the problem (see Sellars 1975, Tye 1984, and D'Ambrosio 2019). Perhaps the easiest solution is to appeal to results from Quine 1960*b* concerning the predicate-functor languages like those in Strawson 1959. But talking in terms of complex predicate functors or quantifying over events or the like either smuggles in object talk through the back door or else assumes a strange kind of metaphysical holism — see Turner 2011 and Filcheva forthcoming. And even if all of these difficulties can be handled, they can't be handled in a simple and natural way. So even if adverbialists can gain expressive parity, we should prefer the sense-data language on practical grounds.

V. The Metaphysics of Sense-Data

The obvious metaphysical question about sense-data concerns the *properties* of sense-data. Is a seeming-red-apple actually red? Most historical sense-data theorists said "yes". In fact, most sense-data theorist accepted the inference from:

(2**) It seems to me that I see a red thing

To:

(3**) I am directly visually aware of something that is red

This move is widely associated with sense-data and nearly as widely decried. Chisholm even called this sort of thing the "sense-datum fallacy".[59]

This inference differs from the inference between (2) and (3*) above. Plugging into the earlier schema, (2**) would only lead to:

(3***) I am directly visually aware of a seeming-red-thing

The question then is whether (3***) entails (3**). More generally: does seeming-*F* entail being *F*, at least for sensible properties like colors and shapes? For a hardcore metaphysician, things get interesting at precisely this point. This certainly seems like a deep metaphysical puzzle about the nature of these important new objects. But a metaphysical approach to this question is unfruitful.

We have here not a deep metaphysical puzzle, but instead a practical question of language engineering — in constructing a sense-data language, should we build in this inference or not? If we do, we might end up saying that a seeming-red-apple sense-datum is red *and* that a red apple is also red. But this sounds a bit odd. To avoid the oddity, we might instead prefer to say that apples are never red, but the sense-data that are caused by

[59] In Chisholm 1957.

some apples sometimes are.[60] Or we might say that there is only one color property here, but that sense-data instantiate properties in a different way than do ordinary things.[61] Or perhaps we could say that while some apples are really red, sense-data are only red* or phenomenally red or p-red.[62]

As this shows, there are a number of options for sense-data theorists. Here I will adopt a version of the final option. I think this is the clearest way of talking, at least in this context.[63] Accordingly, let's distinguish color predicates as applied to sense-data from color predicates as applied to ordinary objects. As far as possible, we can do this while remaining neutral about whether these are two distinct properties. The move I make here for color I also make for the other sensible properties of sense-data.[64] This approach to the properties of sense-data helps to defend sense-data against one of the most powerful arguments against them — Huemer's location argument.[65]

In brief, the argument goes like this. If sense-data have spatial properties like sizes and shapes, it seems that they must also have spatial *locations*. So where are the sense-data located? Where is the seeming-red-apple I'm directly visually aware of? Huemer argues that there is no satisfying answer to this question and so we should reject the existence of sense-data. For example, suppose we try saying that sense-data are located inside our heads. Then there is a red, apple-shaped object *inside of my skull* whenever I am directly visually aware of a seeming-red-apple. That's absurd. But there are also problems with locating sense-data where they appear to us to be.[66] And with locating sense-data where the distal objects are

---

[60] Jackson 1977 and Robinson 1994 accept this kind of view.
[61] Mally said similar things about abstract objects, see Linsky 2014
[62] Even some supposed direct realists, like Searle (2015), appeal to something like p-red in their accounts of color experience.
[63] As already noted in passing, natural language does talk of some kinds of sense-data, for example, after-images. And we say things like "the after-image is red", nothing in my approach is incompatible with this.
[64] This isn't forced — we might allow that shape properties are shared by sense-data and ordinary objects but not color or size properties. If this option is pursued, the location argument might still be avoided by pointing out that some mathematical objects also have shapes without having spatial locations (see below).
[65] Developed in chapter VII of Huemer 2001.
[66] This is Jackson's (1977) view.

located.[67] Both of these options have sense-data out in the physical world overlapping and co-locating with ordinary physical objects. What is in the room? Not just a chair, but also a seeming-chair sense-datum. All of these options for locating sense-data in space are awkward.

There is nothing like a refutation here; the bullet can be bitten. Still, the awkwardness should be avoid it. A response that does so leans on the metaphysical strategy from above (though it is also the best option for sense-data theorists who do not distinguish red* from red). If we deny that sense-data have shapes and sizes and colors in the *same sense* that physical objects do, the pressure to provide sense-data with locations dissipates. The obvious thing to say is instead that (most) sense-data do not have spatial locations. In this, they are similar to numbers and to fictional characters and — unsurprisingly — to experiences. That is: we typically don't think that subjective experiences have spatial locations, except in a derivative sense.[68] As it is with experiences, so it is with sense-data.[69] Sense-data do not have spatial locations. We should be no more flummoxed by the question of *where* sense-data are than we are by the question of *how much they weigh*.

Huemer objects that many sense-data have sizes and shapes, and that this requires them to have spatial locations too. But I have already suggested that sense-data might only have sizes* and shapes*, so the premise of this counter-argument can be denied by sense-data theorists. Further, even if some sense-data have (for example) shapes, so do triangles and other abstract mathematical objects, yet these also lack spatial locations. Huemer also argues that in perception we seem to see things with locations, so if we see sense-data, we

---

[67] Moore (1918-9) flirted with this approach.

[68] Huemer himself admits that experiences don't have locations, but he thinks that saying this requires dualism (I disagree). Somewhat related thoughts are expressed in chapter 4 of McGinn 1999.

[69] Another option that Huemer discusses is locating sense-data in a private, phenomenal space. Something like this is endorsed in Broad 1925, Russell 1927, and Price 1932. But this is not actually an answer to the original question, so I don't think it should be treated as a distinct option. Instead, this idea can be combined with one of the other options. It best coheres with the view that sense-data have no spatial locations. The p-space option might be especially helpful in assigning motion properties to sense-data. There are visual illusions where it seems like there is movement in an objectively static image, and the illusion persists even after you learn that the image is actually static.

should ascribe them locations. But sense-data theorists can say that we do see things with locations – ordinary objects. They have locations, they seem to have locations, but the sense-data we are directly visually aware of do not. The sense-data language as I have developed it here avoids the problems that Huemer raises for the "no location" option.[70]

All of this is well and good for visual sense-data that, while not representational themselves, are necessary for any visual representation of the external world (see section 8). And the same thing also goes for the sense-data associated with the other canonical five senses. Yet today we also recognize additional *internal* senses — of pain, hunger, bodily position, and more. When I have a stomach ache, I am directly aware of a sense-datum. But isn't the ache *in my stomach*? And isn't *that* a location in the world at any given time? If so then these sense-data, at least, must have spatial locations.

I take these as spoils for the victorious army. We may well want to argue that these internal senses generate bodily sensations but not sense-data. But as an answer to the challenge, that seems a bit arbitrary. Personally I would be happy enough to say that the stomach ache sense-datum isn't located anywhere. But I also think that for internal senses we can unproblematically locate their sense-data in the relevant bodily location, thus providing them with egocentric spatial locations. Unlike with a visual sense-datum, in this kind of case, when we locate the sense-datum somewhere in the world, we don't face a problem of overpopulation. It is strange to think the seeming-red-apple sense-datum is located in the same place in the world as the actually red, actual apple. But double-counting oddities don't arise with anything like the same force when I locate my stomach ache *in my stomach*. Some remaining cases might give us pause – is the amputee's pain in their (currently nonexistent) left leg? – but hard cases make bad law.

A related metaphysical question is: what are sense-data made out of? Is there some kind of phenomenal *stuff* out of which they are composed? I think this is analogous to asking

[70] A potential problem Huemer doesn't raise concerns temporal properties – it is natural to say that sense-data have temporal locations, but then physical reality seems to require that they have *space*time locations. I think this move can be resisted in a few different ways, but I also think the premise can be reasonably denied. Any event of directly accessing a sense-datum might have (or be) a spacetime location, but that doesn't mean that the sense-data so accessed also have locations.

what the number 3 is made out of. If these questions have any answer at all, the answer is *nothing*. Sense-data are not made out of any stuff, any more than numbers are. As I will discuss below (section 8) a sense-datum can be *composed* out of other sense-data, but that is different than being *made* out of stuff. The oddity is that sense-data, unlike numbers, are not supposed to be abstract objects. Yet if they have no locations, and aren't made out of anything, how could they be anything but abstract objects? The answer is that there isn't a simple, one-dimensional binary division between abstract and concrete objects. Sense-data should not be classified with ordinary physical objects, but they shouldn't be classified with pure abstract objects either.

More could be said about the metaphysics of sense-data. But let's now turn to the epistemology of sense-data, which is the more historically important issue.

VI. The Epistemology of Sense-Data

The natural epistemological question about sense-data concerns our *access* to sense-data and the *epistemic status* of our beliefs about sense-data. How do we know that sense-data have the properties they have?

The traditional view is that it is literally impossible to be wrong about the properties of your own sense-data. Yet without some nuances added, this is probably too strong. Even prominent sense-data epistemologists like Ayer admitted that merely *verbal* mistakes about sense-data are possible.[71] The real question is whether *factual* mistakes are also possible. The seeming-red-apple seems red, to me. Following the discussion in the previous section, we can say that it is red*. This is a matter of fact. Can I be mistaken about it? Some opponents of sense-data made much of this possibility.[72] Ayer himself later clarified that his position requires only that in matters of sense-data the subject is the ultimate authority.[73]

This is very easy to misunderstand. The claim is *not* that whenever you are directly accessing a sense-datum, you are thereby making an explicit judgement that the sense-

[71] In Ayer 1940.
[72] See Austin 1962.
[73] See the discussion of Austin's 14th argument in Ayer 1967.

datum is *F*.[74] Animals and children (for example) have direct access to sense-data without making any explicit judgements about the properties of sense-data. The point is not about judgement at all, even implicit or tacit judgement, but instead about our non-linguistic *acquaintance* with sense-data.

This is related to an important background issue: the connection between sense-data and dualism in the philosophy of mind. The sense-data theory is sometimes associated with mind/body dualism. This is because sense-data were sometimes stipulated to be non-physical. This association between sense-data and, if not dualism, at least a mysterious concept of consciousness sometimes derives from sense-data theorists themselves.[75] Given this tight association, direct arguments against dualism or p-consciousness have sometimes been taken as indirect arguments against sense-data. This is a mistake. Far from being at odds with sense-data, alternative, access-based approaches to consciousness cohere neatly with reasonable views about both the metaphysics and epistemology of sense-data.[76]

Despite what is sometimes claimed by opponents, sense-data theorists don't need to make "the given" metaphysically basic, determinate, and mysterious.[77] In fact, as I have already been stressing, sense-data is *not* such a given *and was never thought to be*. Sense-data theorists can and do admit that cognitive processing of all sorts goes into generating sensory experience. This point is most naturally made using an access-based approach to consciousness, but even dualist sense-data theorists always take care to make room for it. Witness Robinson:

> It seems to be a phenomenological fact that attitudes, beliefs and anticipations can enter into the structure and tone of the basic phenomenal field, without being fundamentally phenomenal themselves.[78]

---

[74] See the clarifying discussion in Ayer 1945.
[75] See Robinson 1994, for example.
[76] See Block 1995 for the distinction between access and phenomenal consciousness.
[77] The derisive term "the given" is from Sellars 1997. Taylor 1979 mistakenly criticizes sense-data theories for requiring an unconceptualized given.
[78] Robinson 1994, page 207.

Similar points were even made by Locke, in his idea framework.[79] The sense-data framework is independent of our metaphysical theory of conscious experience. The sense-data language begins from the facts of conscious awareness, whatever they may be. And sense-data are not, and were never thought to be, the pure unprocessed input taken in by your senses. Put metaphorically, sense-data is what you get at the sensory finish line, not the sensory starting gate.

This perspective helps to flesh out Ayer's claim that the subject is the ultimate authority about their own sense-data. Here access-based theories of consciousness are particularly well-placed, since they already make the access of a subject metaphysically constitutive of conscious experience. A metaphysical role for access easily gives way to an epistemological role for access. The properties of a sense-datum may well be *partly* or even, in some cases, *fully* determined by the access of the subject. What is crucial, metaphysically, is that a lot goes on *behind* our access, so we need not think of the sense-data as an entirely unprocessed given. What is crucial epistemologically is that our access to the properties of a sense-datum is often much more direct and secure than our access to the properties of ordinary objects. This point, *in some form*, is all that is required for sense-data to play their special role in the early modern project.[80]

Doubts about the existence of the seeming-red patch I am directly visually aware of, in this exact moment, are otiose. Likewise for doubts that the seeming-red patch is actually red*. Such doubts simply cannot gain a foothold. Unlike other skeptical doubts, no plausible possibility of error has here been described. No cogent possibility of error about these matters *can be* described. This contrasts starkly with doubts about the existence and properties of the red apple I currently seem to be seeing. There are many coherent ways for me to fall into error about apples. Not so with sense-data.

---

[79] See Mackie 1976 for some discussion of this.

[80] This might be called into question by Wittgenstein's (1956) notorious "private language" argument. I won't discuss this directly, since I doubt there is a serious challenge to sense-data along these lines. For direct replies to the private language argument by sense-data theorists, see Ayer 1954 and chapter IV of Robinson 1994.

A nice illustration of this difference leans on idealization. If we idealize our ordinary computational and cognitive powers to their limits, all of the errors about sense-data go away, but many errors about ordinary objects remain. If you had no limits of memory, concepts, attention, or computing power, you would be a first-person sense-data oracle. If this idealized version of you had no answer about whether a given sense-datum was *F*, then it could only be because there is no fact of the matter about whether the given sense-datum is *F*. In other words, all of our errors about sense-data fall on the performance side of Chomsky's famous performance/competence division.[81] This illuminates Ayer's claim that, concerning sense-data, the subject is the ultimate authority. And it does so in a way that starkly distinguishes the epistemology of sense-data from the epistemology of ordinary objects.

Even leaving idealizations aside, our direct acquaintance with sense-data gives us epistemological security with respect to their existence and basic qualitative properties. But this epistemological security does not rule out indeterminacy or insecurity concerning more complex properties of sense-data.

To illustrate, let's consider the famous problem of the speckled hen.[82] I seem to see a hen with many speckles. So I see a seeming-many-speckled-hen. How many speckles does it have? I am unable to give a precise answer. This apparently undermines the claim that we have direct access to all properties of our sense-data. In response, Ayer suggested that there is no determinate answer to the question of how many speckles are on the seeming-many-speckled-hen, no fact of the matter. This made his critics uncomfortable, but I think it's the right thing to say in *some* cases. It also dovetails nicely with the indeterminacies standardly posited by access-theories of consciousness.[83]

This is not to deny that sometimes a sense-datum has a non-basic or non-qualitative property and because our access to the datum is fleeting and memory is fallible, our beliefs about whether said sense-datum had said property can be in error. This *can* happen. It *does*

---

[81] From Chomsky 1965.
[82] From Ayer 1940, where it is attributed to Gilbert Ryle; see Chisholm 1942.
[83] Both Dennett 1991 and Warren 2021*a* defend the idea that conscious experience can be *indeterminate* in several ways.

happen. Usually it happens when the properties are holistic and require changes of attention over a given sense-datum. Similar issues arise with mixed properties that have a computational element — does this speckled-hen sense-datum have a prime number of speckles? With non-qualitative properties, or computationally complex qualitative properties, errors are possible. The sense-data language approach is flexible enough to accommodate all of these cases.

A related challenge concerns the intransitivity of relations like "looks the same as". As vagueness and other well-known phenomena demonstrate, α might look the same as β, and β might look the same as γ, while α does *not* look the same as γ.[84] The natural sense-data account of this posits ambiguity.[85] The crucial background point is that it is impossible to directly access all of α, β, and γ at the very same time while still making these same judgements. Given this, it is best to see the tokens of "β" above as picking out two distinct sense-data — β_1 and β_2.

A similar move is also natural to make when considering duck-rabbits and necker cubes.[86] Likewise for some visual illusions. In the checker shadow illusion, one checker square looks darker than another checker square even though the squares are, objectively speaking, the very same shade. Are we here mistaken about the qualitative features of a sense-datum? No. Instead, the sense-data *are* different, but the actual squares in the image are not. There is no contradiction. The same stimulus at or over time can result in distinct sense-data. Once again: sense-data come at the end of a long process of unconscious processing, they should not be identified with either the objective stimulus, nor with an unprocessed sensory given.

When properly understood, the special directness of our cognitive access to sense-data is not undermined by the traditional challenge cases. The same also goes for change blindness and all other related cases discovered and explored by scientists long after the heyday of sense-data. In some of these cases the properties of a sense-datum are

[84] This has been raised as a problem for sense-data by Armstrong (1968).
[85] See chapter VIII of Robinson 1994.
[86] The duck-rabbit is from Wittgenstein 1953. Ayer makes the move I suggest on page 86 of his 1985.

indeterminate. In others, the properties are determinate but are not directly qualitative, so that attributing them places additional cognitive demands on a subject beyond their direct access to the sense-datum. This opens up room for mistakes. In still others, what was initially thought of as a single sense-datum was actually two distinct sense-data. None of this undermines the special epistemological role of sense-data. Doubts about the existence and basic sensory properties of sense-data have a different status than doubts about the existence and basic properties of ordinary objects.

VII. Basic Sense-Data

Above I set-up "the" sense-data language using terms for sense-data that are parasitic on natural language terms for ordinary objects. But in actual fact, we obviously learn the expression "red apple" long before learning anything like "seeming-red-apple". Nothing I have said is in tension with this.

I introduced the sense-data language in the way I did because it makes it very easy to understand. However, a complex sense-datum like a seeming-red-apple is actually *composed*, in some fashion, out of more fundamental sense-data. This applies across all sensory modalities, not just vision. I can be directly aurally aware of a Brahm's-4th-symphony sense-datum. Yet this sense-datum is extremely complex, being itself composed of an arrangement of more basic aural sense-data.

Some philosophers have introduced frameworks for *basic* sense-data. Most famously and influentially, both Carnap and Goodman did this.[87] I won't do the same thing here, but I do think that a basic sense-data language is *possible*. Possible for each individual sensory modality, and also possible for all of them at once. An adequate basic sense-data language would need to include enough basic sense-data properties to ensure that any non-basic sense-datum could be composed out of instances and co-instances of said properties.

If you consider again the Brahm's-4th-symphony sense-datum, you'll get a sense of how complex this might end up being. We would need names for all of the fundamental sounds that together, in a specific arrangement, constitute this complex sense-datum. However, the

[87] See Carnap 1928 and Goodman 1951. See also Peacocke 1983.

difficulties about this do not seem to be "in principle" difficulties. They are instead *practical* difficulties of tedium and unnaturalness. Aiming for basic sense-data, in any modality, is not natural to us in our thinking, talking, or philosophizing. The task involves only the sounds themselves, without care for meaning, emotional shading, context, or anything else. The final output of this painful and tedious process is likely to be quite ugly — a sound of type-C , form-(*iv*) is happening now! — yet ugliness is not impossibility.

I don't think that the in-principle possibility of a basic sense-data language really requires an argument. Never-the-less, here is (a sketch of) an argument. Our sensory system involves several distinct and non-overlapping modalities. These modalities correspond to sub-systems that can each be in various distinct states, all accessible to consciousness. Our experience informs us that these states are structured into local and repeatable qualitative features.[88] That is, our own experiences make this apparent to us, from the first-person. And from the third-person, our hard-won empirical understanding of images, sounds, tastes, smells, and more, reinforces this. Any language that is able to talk about all of these basic sensory property occurrences — red* and circular* or the like — along with all collections and combinations of the same, would be sufficient as a basic sense-data language.

The easiest way to understand composition here is by allowing that any mereological fusion of basic sense-data is itself another sense-datum in the same modality.[89] But appealing to mereology isn't our only option, of course. In fact, our options for individuating sense-data directly parallel our options for individuating ordinary objects, though as discussed in section 5, sense-data are not made of *stuff*. And I stress again that my language-oriented approach does *not* imply that sense-data are created by language. They are not. What I am discussing now is the possible construction of a basic sense-data language, and this too is about how to talk. Different ways of talking will give rise to similar but distinct *concepts* of

[88] Uriah Kriegel has noted to me that some gestalt psychologists have argued that while atomic elements can be distinguished in thought as abstractions, that the *reality* of conscious experience is holistic. I think that this position ultimately requires a Sider-style (2011) rejection of strong quantifier variance, but I won't go into the issue in detail here.
[89] Mates (1981) suggests something like this.

basic sense-data, just as different ways of talking can give rise to similar but distinct *concepts* of ordinary objects. All of this is again backed up by quantifier variance (section 4).

It is useful to talk about sensory runs consisting of all of your sense-data, from every sense-modality, over a given stretch of time. A sensory run usually includes various types of sense-data. The sensory features of any sensory run are themselves factual features of reality. They are not created or generated by language. And an adequate language for talking about experience should be able to describe *any possible* sensory run. Any adequate basic sense-data language will give us this ability. If there are two or more distinct languages adequate for all sensory runs, they will be fully intertranslateable.

In this, an adequate basic sense-data language expressively advances over standard English. Many — *many* — possible sensory runs cannot be described in standard English unaugmented by sense-data talk. You can say that it seems to you that there is a red apple on the table, but it is much harder to describe the sensory experience with any precision. And describing disordered experiences, even in coarse-grained terms, may be impossible.[90] When having an experience that has no order at the level of ordinary object talk, there is no way to talk about the experience by saying it is "as if" you are in some coherent perceptual situation. All such disordered experiences can, in contrast, be referred to and generalized over and talked about in an adequate basic sense-data language. This means that such a language — call it SD-English+ — has more expressive power than standard English.

We don't often feel this lack in English, because in reality, *English is already a sense-data language*. When we speak, in English, about the red afterimage, we are talking about a sense-datum. Likewise for many cases where we talk about sounds, smells, tactile sensations, tastes, and images. Ordinary English doesn't include the term "sense-data", but it includes many *ad hoc* terms and expressions for talking about sense-data. And the same is true of every other natural language too. This, together with the open-endedness of actual

[90] I think that most direct realist theories, whether intentionalist (Huemer 2001, Searle 2015) or disjunctivist (McDowell 1982, Martin 2002) eschew the project of characterizing qualitative experience. From the current perspective, this means they have opted for an expressively impoverished language.

languages, means that actual English probably isn't expressively impoverished relative to SD-English+. Yet this is no comfort to opponents of sense-data, since expressive equivalence comes only because English and other natural languages already include sense-data under different names.

Given all of this, SD-English+ and “English” — English minus its devices for talking about sense-data — provide an example of modest but not strong quantifier variance. Thought the ultimate correctness of this claim depends on the exact equivalence conditions that suffice for languages to be of equal metaphysical merit. On every approach to equivalence aligned with quantifier variance, it will be agreed that many sensory runs can be described in SD-English+ that cannot be described in “English”, and so the languages are inequivalent. SD-English+ is simply a more expressive language.

The general fundamentality of sense-data talk was a commonly held view among early sense-data theorists. One reading of this is modal. When you see a table, you must be directly accessing some sense-data. The opposite is certainly not the case, and this asymmetry is important. Witness Ayer's comments from the mid-nineteen forties:

> For whereas in every case in which it is possible to apply the physical-object language, it is also possible, at least in principle, to apply the sense-datum language, one can conceive an order of experience to which the sense-datum language would have application, but the physical-object language would not. Thus, while it is convenient, for purposes of exposition, to introduce people to the sense-datum terminology by setting forth sentences which refer to sense-data as translations of sentences which refer to physical objects, it would be a mistake to conclude from this that the sense-datum language was nothing more than a technical substitute for the other. There is, on the contrary, an asymmetry between the two languages which may be described by saying that the sense-datum language is logically prior; and this is shown by the fact that, while referring to sense-data is not necessarily a way of

> referring to physical objects, referring to physical objects is necessarily a way of referring to sense-data.[91]

I agree with almost all of this. My only disagreement is with the last sentence. When the paper from which this quote was drawn was written, Ayer still accepted phenomenalism. In later work, like virtually every other historical phenomenalist, Ayer rejected phenomenalism as unworkable.[92]

Certain central sentences, expressible in the sense-data language, are epistemologically special. They do not go beyond what is directly accessible to consciousness in a moment of awareness. We might call these "basic experience claims".[93] Many of them can be made in SD-English, and *all of them* can be made in SD-English+. This is not a merely formal claim, since it is not about *labels*.[94] The flux of experience is real, it is the empirical data available to consciousness. And the sense-data language is the best way of talking and thinking about the data of experience.

This is a way in which the sense-data language is "logically prior" to the ordinary object language. Another way, discussed in the next section, concerns the role that sense-data plays in our philosophical theorizing. Yet another, discussed thereafter, concerns the epistemological role of sense-data in facing down the unavoidable challenge of skepticism.

VIII. Sense-Data Theories of Perception and Color

Consider again the three illustrative claims from section 3:

(1) I see a red apple on the counter in front of me

(2) It seems to me that I see a red apple on the counter in front of me

---

[91] Page 312, Ayer 1945. In the context, Ayer was responding to and agreeing with claims from Price 1941 about Ayer 1940.
[92] Ayer became a realist of a sort, see the extended discussions in his 1968 and 1973.
[93] Similar terminology is used in (for example) Russell 1940 and Ayer 1968.
[94] See Ayer 1940 on the errors of formalism for important clarity about this.

(3*) I am directly visually aware of a seeming-red-apple-on-the-counter-in-front-of-me

There is a natural explanatory order on these claims: (2) is at least partly explained by (3*) and (1) is partly explained by (2) and partly explained by the fact that there is a red apple on the counter in front of me that I stand in a relevant causal relationship to.

Why didn't I say that (2) is "fully" explained by (3*)? Because there is an important ambiguity already touched on: when I say it seems like *F* to me, often, though not always, some conceptual ability is imputed to me with respect to *F*. But as also stressed throughout, sense-data themselves are not inherently intentional or representational. So I sometimes prefer to say that (2) is *partly* explained by (3*) and *partly* explained by some fact about my concepts and how I am exercising them in this given situation. If we instead build this conceptual understanding into the truth conditions of (3*), then we can say that (3*) itself fully explains (2). Likewise if (2) is used without ascribing conceptual understanding to the subject (which it often is).

Don't be misled by any of this. Direct realists make much of the intentionality of perceptual experience.[95] Sense-data theorists can agree with them, at least at the level of (2). Yet, when so understood, (2) is not metaphysically or epistemologically fundamental. But a sophisticated representational theory of perception can be combined with a sense-data theory. This is because sense-data are the data of conscious sensory experience, but this data is itself built out of information taken in by the senses and processed and altered in various ways on its way to the global workspace. The representational role is, in an externalist sense, a generalization of the cognitive role played by the deliverances of perception.[96]

Sense-data theorists have flexibility in offering theories of both perception and perceptual properties. To illustrate this, focus on color.[97] Above, in section 5, I distinguished between

[95] Both Huemer 2001 and Searle 2015 stress this.
[96] See Cao and Warren forthcoming.
[97] For an overview of color for philosophers, see Hardin 1988.

red as a predicate of ordinary objects and red* as a predicate of sense-data. Even supposing that these are two distinct properties, they are obviously very closely related. In fact, the ordinary object property of red can be explained in terms of the sense-data property of red*.

An appealing approach for this is a sense-data version of color dispositionalism. The key observation is that when I see a red object, in normal conditions, I am directly visually acquainted with a red* sense-datum:

> **SD Color Dispositionalism.** An ordinary object is red just in case the object would cause ordinary human observers in normal conditions to be directly visually aware of a red* sense-datum appropriately associated with the object

This is an explanatory principle where the right-hand-side explains the left-hand-side.

This simple color dispositionalism would need to be elaborated and refined in various ways for a full account. This would require explaining the relevant notions of "standard observer" and "normal conditions", though I don't think that either task poses impossible problems. The final bit, about the red* sense-datum being "appropriately associated with the object" is meant to rule out a few lingering nonstandard cases. An object that produces a blinding flash of light might cause observers to see a red afterimage. An object that sends out nanobots to tinker with the red*-producing center of an observer's brain might cause the observer to see red rockets. Intuitively, neither of these objects is red, despite these dispositional features.

Color dispositionalism of this kind goes back, at least, to Locke. Modern forms of color dispositionalism have also been defended in the contemporary literature.[98] But the modern forms typically eschew sense-data. This omission threatens dispositionalism with circularity. Put baldly then, suppose that *being red* is explained in terms of *looking red*. But what is looking red? Don't we need some independent account of redness to avoid circularity? There may be other ways out, but the circularity challenge doesn't even arise for sense-data

[98] For examples and discussion see Bennett 1971, Dummett 1979, Hardin 1988, and Johnston 1992.

versions of color dispositionalism. The sense-data approach doesn't explain red-facts in terms of other red-facts, it instead explains red-facts in terms of red*-facts. What goes for color can also go for all of the other so-called "secondary" qualities of objects. Locke smiles down on us from philosophical heaven.

I won't go into further detail here. Obviously these remarks don't amount to a worked-out theory. My goal was only to suggest that the sense-data language can be leveraged to give plausible theories of perception and secondary properties. Given this, sense-data theorists are not boxed in when it comes to theorizing about perception, the mind, and the relationship between the mind and the world. However, the *ultimate* question about this relationship is the problem of external world skepticism. Echoing Reid's similar claims about the idea theory, *many* philosophers have claimed that if you accept the sense-data theory you are forced into skepticism. So since we shouldn't be skeptics, by *modus tollens*, we shouldn't be sense-data theorists either. This challenge is the final boss for sense-data proponents. It's time to face it.

IX. Facing Up to the Problem of Skepticism

The challenge of early modern skepticism was to vindicate our beliefs about the external world on the basis of our ideas alone. The modern version of this challenge is to vindicate our beliefs about the external world on the basis of our sense-data alone. Even posing this challenge properly requires something like the sense-data language. If this challenge is, once posed, unanswerable, that is not itself a reason for rejecting sense-data. The sense-data language isn't *creating* the skeptical problem, it is merely *enabling* its clear statement. Avoiding the sense-data language because you fear an unanswerable challenge is the intellectual equivalent to sticking your fingers into your ears and humming.

Many contemporary discussions of skepticism leave the traditional problem completely untouched. They aim instead to show how we can know that the cat is on the mat despite not knowing that we're not dreaming.[99] Or that the burden of proof is on the skeptic, so we

[99] See Dretske 1970.

can wave our hands, one at a time, and be done.[100] Or that in most contexts, the skeptical scenarios can be appropriately ignored.[101] Or that because of semantic externalism we can't even think the thought that we might be in a skeptical scenario.[102] Or that ordinary knowledge is possible because, given that the external world is as we think it is, our beliefs about cats and hands couldn't easily be mistaken.[103] Or that the nature of belief itself makes accepting a global skeptical position impossible.[104] Or that it seems like an external world exists and this seeming can be taken at face-value.[105] Or…well, you get the point.

With these labors has philosophy at last solved the great skeptical problem? Let's perform a small thought experiment. Suppose we resurrect Descartes and the other early modern giants. Suppose also that we teach them today's English and allow them to study contemporary works on skepticism.

> Question: Would they be impressed by the progress we've made on this most difficult of philosophical problems?
>
> Answer: No. *Not at all*.

At least, this is what I think would happen if Descartes and Locke and Berkeley and Hume and friends were faced with the array of options just listed off. It's also what happens to many undergraduate students every year. Having been gripped by the skeptical problem, the student is invariably disappointed at the responses. Many of these students are eventually brow-beaten into holding their tongues when faced with various types of externalism and a shifting burden of proof, but they never shake the feeling that all of this misses the point. This is a rational response on their part.

---

[100] See Moore 1939.
[101] See Lewis 1996.
[102] See Putnam 1981.
[103] See Williamson 2000.
[104] See Greco 2012.
[105] See Pryor 2000.

I am not saying that all contemporary work on skepticism is worthless. I am not even saying that any of the work mentioned above is worthless. I am merely saying that, even if these works succeed perfectly at their aims, they are not relevant to the traditional problem of skepticism posed by Descartes.[106]

My point isn't merely that the skeptical challenge should be understood in an *internalist* sense rather than an *externalist* sense. There is more to it than that. Most contemporary internalist responses to skepticism don't resolve the real problem either. Let me illustrate this by considering the last view mentioned in the litany above.

Many so-called "internalists" accept something like the following principle about *prima facie* (internalist) justification:[107]

> **Seeming-to-Justification.** If it seems to you that *p*, then you are *prima facie* justified in believing that *p*

This principle is accepted by many direct realists about perception, but sense-data theorists don't need to deny it. The principle has been mooted as a reply to skepticism — it seems to me that there is an external world, so I am *prima facie* justified, even justified in an internalist sense, in believing that there is an external world. Voila.

Fair enough, but this doesn't address the real problem. Our *prima facie* justification can be defeated, and whatever *prima facie* justification "seeming" provides surely *is* defeated when we consider skeptical scenarios. You look to the world and it seems to you that there is a barn in the field. By the above principle, you are *prima facie* justified in believing that there is a barn in the field. Yet this is actually the famous field of fake barns, built by a family of philosophers for the epistemological edification of the public. When I tell you this, your epistemic situation changes. Whatever *prima facie* justification you had for your barn-belief is over-ridden. If you continue to form barn-beliefs here on the basis of perceptual seemings,

---

[106] We might compare Chalmers's (1995) contrast between the "easy" and "hard" problems of consciousness. I originally used parallel terminology here, but ended up thinking it was more distracting than helpful, though I have kept the Chalmers-inspired section title.

[107] For relevant discussion see (for example) Pryor 2000 and Huemer 2001.

you are making a serious mistake. The same is true even if I lied to you, or if I was wrong about our location. Either way, a possibility of error has been made salient and it must be addressed. Addressing these possibilities of error *is* the problem.

When external world skeptical scenarios are presented we find ourselves in an analogous situation. There are many coherently describable scenarios where I have the very same sensory run I just had but without any external world. Given these possibilities of error, how can my external world beliefs possibly be rational? Without an answer to this sort of question, the real skeptical problem is untouched. This is true even if my *prima facie* justification remains. The skeptical questions highlight the gap between *prima facie* and all-things-considered justification.

When I talk about external world skepticism or the real skeptical problem, I am talking about the traditional, early modern version of the problem. An acceptable resolution to this problem cannot go beyond the first-person starting point. We have our sense-data and whatever processing and reasoning can be launched from it, *a priori*, virtually nothing else. This isn't meant as a controversial philosophical claim, but instead as something that is obvious once it is pointed out. We face the world looking out from our eyes, not from God's point of view. If someone asks you to back-up your fundamental beliefs in the external world without assuming anything about said external world up-front, you owe them an answer. Descartes has asked us, so we owe him an answer.

At least, you owe an answer *if* the first-person perspective is coherent. If it is not, then the challenge itself is incoherent. The famous ordinary language criticisms of sense-data would have shown this, if successful. My defense of the sense-data language in the first eight sections of the paper is my reply to their attempt. Yet, the failure of the ordinary language challenge was already obvious even without a detailed defense of sense-data. Mates provides a lovely statement of this point:

> To the complaint that the epistemologists who have introduced the term "sense data" and its various associates have not succeeded in making us aware of what they are talking about, the response is a simple denial. In fact, to complain that a typical sense-datum epistemologist like G.E. Moore, with his endless explanations … is

> careless about the use of language is nothing short of preposterous. … Often [the critics of sense-data] tell us authoritatively that "You *can't* say this" and "You *can't* say that," when the this and the that are precisely what large numbers of native speakers *do* say, managing without any difficulty to communicate with one another thereby. Thus, "I can perceive only my own perceptions" may sound linguistically odd at first hearing, and one might be inclined to retort, "No, you perceive books and chairs and trees and the sky, but you don't perceive *perceptions*." But after the epistemologist gives all his explanations of the word "perception," the proposition appears as a truism. It's too late then for the Ordinary Language Philosopher to tell us, "But you *can't say* "I am perceiving a perception."" We *can* say it; we *do* say it; and it seems obviously true.[108]

Indeed. And the quantifier variance analysis of the sense-data language theoretically vindicates this intuitive reply.

The sense-data language is coherent. You understand exactly what is meant when Descartes, Locke, Berkeley, Hume, Russell, Ayer, and I say things like, "You are not ever directly acquainted with ordinary objects, only with sense-data". You understand it and you know that it's true. Trivially and obviously and blamelessly true. From here, we stare into the mouth of the dragon. The problem – the *real* problem – of external world skepticism is now upon us. It is one of the deepest and most disturbing of all philosophical problems. You are free to avoid answering, but don't pretend that by ignoring the problem you have *thereby* answered it.

X. Patterns in Experience

A sensory run consists of all of your sense-data over an unbroken temporal interval. All of the sense-data from all of your senses over that interval. You are directly acquainted with this data. Yet there is more going on in your head than just this. The first-person perspective must also allow acquaintance with (at least) emotions and the activities of the will.

[108] Mates 1981, pages 152-153. See also Mates 1967.

It is also important to recognize that direct awareness takes time. We shouldn't think of acts of awareness as being instantaneous. Instead, we should recognize what Williams James called, following some precursors, "the specious present".[109] Some early sense-data theorists already recognized the importance of this.[110] It secures the passage of time as something we have direct awareness of, from the inside. Whenever I talk about direct awareness "in a moment" or "at a time", it should be understood in this way.

There is room for debate about exactly what is first-person accessible to consciousness.[111] As an illustrative model, I will limit our first-person access in a moment to the following:

> sense-data from all sensory modalities (in full generality this includes data from interior senses too)
>
> emotional states (understood broadly)
>
> activities of the will (including internal willings and inclinations)
>
> temporal awareness (of the specious present and local temporal change)

I tentatively think we can understand imagining, remembering, and all other cognitive phenomenology using only these components in various combinations. Others might prefer to understand willing itself in a sensory or emotional way. Still other variations are possible. Most ways of cutting up this pie would serve for my purposes, provided that *sui generis* states of rational insight or intellectual intuition are *not* included.

I will say that these accessible states of the sensory system, the affective system, and the will (the first three) at a moment of time (the fourth) constitute one's accessible (or first-person) mental state at that time. An accessible mental state over an extended interval consists of one's accessible mental states over the moments in the interval. Since the

[109] In James 1893.
[110] For example, see Russell 1914*b* and Broad 1923.
[111] Thanks to Uriah Kriegel here.

"moments" are not literally instantaneous, we should think of them as overlapping and approximately continuous.

This is all that we have to work with. This and whatever we can get *a priori* from this basis. Nothing else. Of course, when discussing the skeptical problem we usually allow as data your accessible mental state over your entire existence up to the present. This is a simplification. Ultimately, this full history is also something that must be secured *in each moment* on the basis of what you are directly acquainted with in that moment.[112]

Your accessible mental state, even over your entire life, does not deductively entail the existence of the external world. This was illustrated forcefully by Descartes. Ever since his discussion, nobody has thought otherwise. But in his *Meditations* Descartes himself tried to solve the problem deductively in a more circuitous way, by first proving the existence of God, *a priori*, and then using God to vouchsafe the evidence of the senses. Today, this strategy is unlikely to appeal even to theologians. The problem is simply that, with apologies to Saint Anselm, God's existence cannot be proven *a priori*. The *a priori* cannot turn water into wine; there is no epistemic magic in the world.[113]

In the previous section I galloped over a broad landscape, claiming that much of the contemporary philosophical work on skepticism is powerless against the real problem. Much, but not all. There is one plausible strategy for solving the real skeptical problem. One and only one. The strategy is to argue that we are *rational* to posit the external world as the *best explanation* of the patterns and regularities in our Cartesian mental states. This broad strategy is known as abductivism or the IBE-strategy. The phrase "inference to the best explanation" was only invented in the nineteen-sixties, but the type of reasoning is much older.[114] The IBE-reply to skepticism, or something like it, was offered, at least implicitly, by many canonical sense-data theorists — Russell, Broad, Ayer, and Jackson, for example.[115]

---

[112] This point is made in, for example, Rinard 2017. Traditional sense-data epistemologists were well aware of it.
[113] For a theory of the *a priori* without magic, see Warren 2022*b*.
[114] The phrase derives from Harman 1965.
[115] See Russell 1912, Broad 1925, Ayer 1973, and Jackson 1977. Russell and Ayer also endorsed phenomenalism at various respective points, though both later flirted with a

The IBE-strategy has more recently been both widely advocated and widely criticized.[116] The strategy is often presented in a somewhat confused form, even by proponents. For instance, perhaps the most prominent version of the strategy, from Vogel, isn't fully explicit about what is being explained or how. Vogel also appeals to (in my view) implausible principles concerning explanation and necessary truth. A proper version of the IBE-strategy must initially appeal only to facts stated in the sense-data language or some equivalent.

The IBE-strategy refines the simple causal strategy implicit in Locke's *Essay Concerning Human Understanding*. That simple causal approach claims that external objects are the cause of (many of) our ideas. It develops this thought by holding that our ideas of objects "resemble" the things in the world that they are ideas *of*. But it is unclear how to justify this claim from the first-person on empiricist grounds. If we can only ever perceive our ideas, then we can never break through the "veil of perception".[117] Thus, we might be justified in claiming that two ideas resemble each other, but never in claiming that an idea resembles something in the world. Both Berkeley and Hume savaged Locke's approach, and Hume infamously put causation itself under the empiricist microscope. It also didn't help that Locke never set up his explanatory target carefully or properly.[118]

Probably the best early discussion of this is from Hume, though he adopted it only as a descriptive story of *why* we believe in an external world — there is no normative element in his account. Like Locke, Hume often contaminated his discussion of the patterns we are trying to explain by building in information about the external world. Even so, his discussion is impressive and historically important. Here is Bennett commenting on this section of Hume's *Treatise*:

> It is extremely difficult, full of mistakes, and—taken as a whole—a total failure; yet its depth and scope and disciplined complexity make it one of the most instructive

broadly IBE-style defense of scientific realism. But Ayer later disavowed this understanding of his approach in favor of a more pragmatic interpretation, see Ayer 1992.

[116] See Slote 1970, Mackie 1976, 1982, Cornman 1980, Vogel 1990, Fumerton 1992, Bonjour 1999, Neta 2004, Beebe 2009, Gifford 2013, and Rinard 2017.

[117] This famous phrase was introduced in Bennett 1971, page 69.

[118] See the discussion in Bennett 1979.

> arguments in modern philosophy. One philosopher might be judged superior to another because he achieved something of which the other was altogether intellectually incapable. By that criterion Hume surpasses Locke and Berkeley—because, and only because, of this one section.[119]

Hume's treatment was also nicely discussed in a now little-remembered book by Price.[120] I might be the only person to have read this book in the twenty-first century, but it remains worth reading.

Hume focused on the constancy and coherence of our impressions (remember that, for him, "ideas" were fainter copies of impressions). Price cleaned up this approach. Sometimes a sensory run consists of a sequence of sense-data — or, to use Hume's term, *impressions* — over time:

$$E_1, E_2, E_3, E_4, E_5$$

A series like this exhibits constancy just in case each member of the series is qualitatively identical to the next member of the series. Coherence is then defined as a relation between two different series, one with a "gap" and one without:

$$A, –, –, –, E$$

$$A, B, C, D, E$$

The first and last terms of each series are qualitatively identical, and this is what coherence consists in. The sense-data in the first and last positions of two distinct sensory runs can be indistinguishable, even if the middle parts are very different. The Humean idea is that we posit stable "impressions" to account for both constancy and coherence. In this way the imagination projects itself *into* perceptual gaps.

Hume seemed to think it was "impressions" *themselves* that were assumed stable and independent. He even sometimes talks as if the very idea of external, mind-independent objects is a bit of nonsense that philosophers invented to gloss the talk of the vulgar. In this, he – and Berkeley before him – might have been overly focused on mistakes by Locke. Mistakes that they also overstated. Later, Reid savagely mocked this strand in empiricist

[119] Page 313, Bennett 1971.
[120] See Price 1940.

thought and rightly exposed it as an absurdity. The vulgar themselves posit external, physical objects. The only invention of philosophers is the claim that they don't.

This is one problem with Hume's account, another is that – as already noted – the patterns he spells out are often mischaracterized in his discussion. His examples repeatedly build-in or implicitly assume the existence of external objects and so aren't actually patterns *in impressions* at all. As both Bennett and Price note, greater care must be taken for a proper treatment.

The sense-data language allows us to take this greater care. Our accessible mental states exhibit many different patterns and regularities. Using the sense-data language we can talk about these patterns in a very *general* way, over our overall accessible mental states, or in *particular* ways, over a single sense-datum at or over time. Here is a non-exhaustive and schematic list of patterns in our first-person mental states:[121]

(1) synchronic structure in the sense-data of each sensory modality — visual experience (for example) is not a random flux or even a kaleidoscope, as I look at the room right now there are clusters and regularities even in the moment

(2) synchronic regularities connecting the sense-data of different sensory modalities — when I have the visual sense-datum of seeing-fingers-snapping, I also have the tactile sense-datum of feeling-fingers-snapping, and the aural sense-datum of hearing-fingers-snapping

(3) diachronic structure in the sense-data of each sensory modality — examples include but aren't limited to Hume's constancy and coherence

[121] Two clarificatory points: (*i*) synchronic patterns still have a small amount of temporal structure, due to the specious present being the temporal unit of psychological awareness; (*ii*) if our data consists only of the specious present, then the diachronic patterns will not be directly given, but rather inferred in the same way as the external world.

(4) diachronic regularities connecting the sense-data of different sensory modalities — when playing the piano over an interval, visual, tactile, and auditory sense-data co-occur in regular and repeated ways

(5) correlations between willing and sense-data in each sensory modality — I can imagine a pink elephant, I can also, less directly, make a sensory run repeat but in reverse order, by walking back into the room, turning my head from side to side, or the like

(6) correlations between sense-data and emotional states — when I have a sense-datum of a seeming-hammer-hitting-my-thumb, an emotional state of extreme pain and distress immediately follows

This list is neither exhaustive nor extremely detailed, but even at this level of description the point is powerful. There are *striking* patterns *over* sense-data both synchronically and diachronically, as well as *striking* patterns *across* the different components of our accessible mental states. Our subjective experience of the world is extremely structured and well-patterned.

The sense-data language enables us to describe and theorize about all of these patterns in the most general way. If you don't have the sense-data language, or something expressively equivalent, you might be able to talk about some of the patterns, but not all of them. Merely saying that you seem to see a brown chair suggests that there is a type-(1) pattern. If you seem to see a brown chair over an interval of time, you can even capture some type-(3) patterns. Yet I would also seem to see a brown chair over an interval of time if my vision was cycling between distinct brown chairs at one second intervals. The actual patterns to be explained are difficult to capture at this level of description, perhaps impossible.

It is also difficult to capture type-(5) patterns with this way of talking. I seem to see a brown chair, I close my eyes, I seem to see a brown chair again. I seem to see a brown chair, I walk out of the room, and I seem to see a brown chair again, with an experience

that is roughly reversed. The general difficulty is that seeming to see an *F* is too coarse-grained. Seeming to see an *F* is compatible with a great many distinct subjective experiences. Many of these are not similar to each other in a qualitative sense. In order to pick out the crucial patterns, you will need to introduce more and more complex descriptions — seeing a brown chair from the side at twenty paces in neon lighting with a head angle of twenty degrees. Complexities of this kind can be added almost without end, but at a high price. If this approach can reach expressive parity with the sense-data language at all, it does so at the cost of extreme unnaturalness.

In stark contrast, the sense-data language makes the statement and description of these patterns both natural and easy at every level of granularity. There is no better or more flexible way to talk about experience than with the sense-data language. But, of course, the patterns are there whether we use the sense-data language or not. We did not make them with language, *we* did not make them at all. The sense-data language allows us to talk and reason and refer to and generalize about these empirical patterns.

Once we have identified these patterns, we see that they call out for an explanation. The best explanation of them is that there is an independent external world that causally impacts upon us in regular ways. This is the answer to external world skepticism. I have stated it baldly, but still, it is the answer. Let me explain.

XI. Explaining the Patterns

There are patterns in our first-person mental states, patterns we can access and identity from the inside. I have said that these patterns call out for an explanation, but why is that?

We could simply take it as an *a priori* principle that certain striking facts need to be explained. Yet something more systematic can also be said. Think of all of the possible qualitative experiences you could be having in a moment or over an interval. To limit the scope of things, let's focus just on vision. Most of the possible arrays of visual sense-data that you could be experiencing at any given moment are utterly incoherent to you. They don't even have the structure of a kaleidoscope experience, let alone the structure of ordinary visual experiences.

The number of possibilities here is — *at least* — finite but astonishingly large, even in a single interval of awareness.[122] This is most obvious for vision, but something similar holds for the array of sense-data that could potentially be delivered by *any* of our senses. Mixing and matching these possibilities gives the number of possible overall sensory states, both at a time and over time. The numbers are truly astronomical. Now add in the states of the will and the emotions, and you have a staggeringly, extra-astronomically large finite number of possible states of conscious experience.

The number of such states is certainly much larger than one billion. But to cleanly illustrate the point, suppose there are a mere one billion distinct overall experiences you could be having in a single second. Let's also suppose that your accessible states develop discretely over time. In reality, development over time is much messier, likely leading to even more possibilities. Since these states are *logically* independent of each other, the number of possible experiential runs over ten seconds of time is:

$$1{,}000{,}000{,}000^{10}$$

If written out in full decimal notation this would be a one followed by *ninety* zeroes. It's almost a googol (not really, but you know what I mean). And remember that this is over only *ten seconds*, and remember too that this estimate drastically understates the actual number of possible experiences even over this brief interval.

Out of this vast number of possible experiences only a very small fraction exhibit any of the regularities mentioned in the previous section. If sensory runs of 10 seconds in duration were chosen at random out of all possible experiences, it would be very, very unlikely that the chosen runs would exhibit any regularities of the kinds mentioned in section 10, let

[122] The sense of "possibility" at issue here is largely *a priori*, nothing about the physical workings of the sensory system is being assumed. In fact, nothing about *physics* is being assumed, at all.

alone all of them. This combinatorial point gives bite to the idea that regularities in experience are unexpected — they have a very, very low *a priori* prior probability.[123]

This reasoning assumes that each *a priori* possibility should be treated on a par. This can be supported by some kind of general symmetry principle for epistemic probability. The most famous of these is the infamous principle of indifference. According to indifference, absent any evidence to act as a symmetry breaker, it is rational to assign equal probability over the cells in a partition of possible outcomes, even *a priori*. Incautious formulations of indifference risk inconsistency, but there are many proposals for avoiding these troubles.

The very idea of *a priori* rationality constraints on belief and degrees of belief *sounds* controversial. As a matter of descriptive fact, it *is* controversial. Yet it shouldn't be. If *a priori* rationality constraints are not accepted, then there is no difference between rationality and irrationality.[124]

Of course, to fully vindicate either IBE or *a priori* probability principles from the first-person starting point requires solving Hume's problem of induction. I hope readers will forgive me for not trying to do that here. Instead I will briefly explain one shape that Hume's challenge takes in this particular context, and why I think it can be answered.[125]

Suppose that I win the lottery by playing my birthday numbers. I might think that the connection between my birthdate and the winning number requires some special explanation. I might think that, but I'd be wrong. There is no explanatory hole left if we take this alignment as a brute coincidence. The feeling to the contrary is an illusion. We all agree about this. Any feeling to the contrary is an illusion. The inductive skeptic thinks the idea that the patterns in experience are unlikely and so need to be explained is also an illusion.

---

[123] My discussion here largely dovetails with the Huemer's (2016) discussion of the brain-in-a-vat case. Huemer isn't advocating an abductivist reply, but his appeal to *a priori* probability is, I think, in line with abductivism.

[124] An early statement of (a version of) this point is on pages 66-68 of Putnam 1971.

[125] See chapter 1 of Ayer 1972 for a relevant, modern defense of Hume. Obviously, the "new riddle of induction" from Goodman 1946, 1955 is also relevant here. See the essays and annotated bibliography in Stalker 1994 for a comprehensive overview of the first fifty years of work on this problem.

Some proportion out of all possible sensory runs are patterned. And any other collection of experiences of the same size is just as probable, *a priori*.[126] But then *whatever* particular experiences we have, patterned or not, we can engage in analogous reasoning. Suppose we have apparently random experiences over a time, where a multicolored quasi-patch forms and instantly vanishes in the upper left of the field of vision. Call this a property *MC*, of experiences. The Humean says that *MC* is just as unlikely *a priori* as the richer patterns discussed in the previous section. This Humean further claims that the same points also go for *any* collection of possible experiences, and any features an experience might have. So the idea that patterned experiences are especially unexpected and call out for an explanation is — this skeptic claims — an illusion.

The Humean's reasoning is tempting, but it is fallacious. No less a luminary than Doctor Manhattan once succumbed to a similar fallacy. On Mars, the good Doctor said to Laurie:

> Thermodynamic miracles … [*sic*] events with odds against so astronomical they're effectively impossible, like oxygen spontaneously becoming gold. I long to observe such a thing. And yet, in each human coupling, a thousand million sperm vie for a single egg. Multiply those odds by countless generations, against the odds of your ancestors being alive; meeting; siring this precise son; that exact daughter … [*sic*] Until your mother loves a man … and of that union, of the thousand million children competing for fertilization, it was you, only you, that emerged. To distill so specific a form from that chaos or improbability, like turning air to gold …[*sic*] That is the crowning unlikelihood. The thermodynamic miracle.[127]

Despite nigh-omniscience, Doctor Manhattan is mistaken. There is a sense in which Laurie's existence is unlikely. As was mine, as was yours. It was extremely improbable for *you* to come into being. Yet this particular improbability can only be specified *post hoc*. That is, it

---

[126] The notion of "size" can't be understood simply, in terms of cardinality, except in the finite case. Handling probability over infinite outcome spaces will require more subtlety and discrimination in our notion of "proportion". This is a problem for everybody though, nothing specific about my set-up.

[127] Quoted from Gibbons and Moore 1987.

is entirely *backwards*-looking. Backwards-looking both at the level of individual coupling and at the level of generation-to-generation transitions. It is true that the particular sperm-egg combination that resulted in your existence was exceedingly unlikely, but that doesn't make your existence miraculous.

Some person was almost certain to result from repeated couplings. *Some ticket* was almost certain to win the lottery given repeated drawings. Be careful about the order of the quantifiers here — in many nearby possible worlds, your parents created a different child, in many nearby possible worlds, a different winning number was drawn. There are no *post hoc* miracles. If a chance process has *n* possible outcomes, for a large *n*, each particular outcome can be astronomically unlikely. Yet the chance process itself might be such that one of these outcomes will almost certainly occur. In this kind of case, it is fallacious to claim *after the fact* that a miracle has occurred. Otherwise, we could be certain in advance that no matter what happens, a miracle will have occurred. This kind of supposed miraculousness is an illusion.

Yet the miraculousness of the patterns in our experiences is not of this kind. It can instead be specified while looking forward. The continuities in this moment of experience are not expected just because some experience must be occurring. And the continuities between the future and the past are not expected given simply that *some* experience will be had in the next second, over the next ten seconds, or over the next year. There is a real difference between the patterns that call out for an explanation, and the "patterns", that do not. The inductive skeptics (and Doctor Manhattan) are wrong here. They are treating distinct epistemic situations as if they were alike.

The patterns in our accessible mental states demand an explanation. The external world hypothesis provides an explanation. It makes the patterns expected with a theory that has all of the explanatory virtues we seek.[128] The external world hypothesis is, far and away, the *best explanation* of the patterns in our first-person mental states. In fact, it does not have a

[128] In terms of probability, this means that the external world hypothesis has a high *a priori* probability and the patterns in experience have a high *a priori* likelihood conditional on the external world hypothesis. For this to work, the external world hypothesis might require that causal laws be understood in a non-Humean fashion.

close rival, so it is rational to accept the hypothesis. More than that, rationality *demands* that we accept the external world hypothesis. This argument can be developed in several different ways. One way is to construe IBE as a rule of inference, and to use as premises the claims that the external world theory meets the conditions for an IBE-inference. Given our empirical evidence, in the form of sense-data and our accessible mental states, it is then rational to infer the existence of the external world by inference to the best explanation.[129]

In a full treatment we would need to distinguish between many distinct external world hypotheses and treat them all separately, appealing to different patterns in our first-person mental states in each case. Without going through each case here, let me simply claim that IBE-reasoning can support, from the armchair, at least the following distinct external world hypotheses:

> *INDEPENDENCE*. There is an external world that is largely independent of me. It does not depend on me for its existence or general nature.
>
> *STABILITY*. The objects in the external world and their features are fairly but not perfectly stable over small intervals of time.
>
> *UNIFORMITY*. The objects in the external world and their features are fairly but not perfectly uniform across the observed and unobserved portions of the world.

[129] Among these conditions: that the external world explanation was sufficiently good and that no other explanation meets this threshold, among other things. For important discussion, see Lipton 1991. Another option is to try to fit IBE reasoning into a Bayesian framework. There are some arguments that this is impossible — see Van Fraassen 1989 and Roche and Sober 2013 — but there are also interesting responses — see Okasha 2000, Lipton 2001, Huemer 2009, and Weisberg 2009. The connection between IBE and Bayesianism is an important topic in epistemology, but I won't go into detail here (but see the previous footnote).

TRACKING. The features of the external world are imperfectly tracked by my sensory system.

BODY. My body is a part of this external world, housing my sensory system and manipulating the body-external environment in a limited way.

BASE. The external world is the base world, in that there is no distinct world in which the world we know through experience is embedded.

If I'm right about all of this, then a sequence of inferences to the best explanation will allow us to rationally believe in an external world that has roughly the features we now take it to have. A world that is not dependent on the mind or the will. A world that isn't a dream or a hallucination or a simulation.[130]

This reasoning doesn't establish any esoteric *metaphysical* claims about the precise nature of the external world. Nor does it defeat *a posteriori* skeptical challenges like Bostrom's simulation argument.[131] Instead, it answers the traditional skeptical problem and shows us why we should not be solipsists or phenomenalist. We should be realists about the physical world, and should carefully build our overall theory of this external world step by cautious step, ever mindful of the connection between our theories of this reality and the direct evidence provided by our sense-data.

Even very deep into this process unexpected results continue to show up. Here is one small example. Scientific investigation of vision revealed the existence of the blind spot. I don't know about you, but I first learned about the blind spots in my eyes by reading about them. I was then able to do self-experiments demonstrating their existence, making little black dots on a page disappear. We have here a case where a pattern in sense-data that

[130] In addition to being backed by radically different arguments, this goes beyond the "structuralist" reply to skepticism from Chalmers 2018. This can be seen by considering skeptical scenarios that aren't ruled out by a structuralist reply, see Berry 2019 for discussion.

[131] From Bostrom 2003, though Bostrom actually only argues for a disjunctive conclusion containing the simulation hypothesis as a disjunct.

does *not* make sense in terms of the existence of the external world, in the normal way, is none-the-less predicated and explained by the external world hypothesis at a suitable stage of its development. In this and a million other ways, patterns in sense-data rationally support more and more refined and expansive external world hypotheses. This is the solution to the real challenge of external world skepticism.

XII. A Plea for Analytic Empiricism

I have argued for a number of related claims. That the early modern starting point — including the crucial sense-data component — is coherent (sections 1-4). That we can adopt a sense-data language when thinking and talking about our experiences and our relation to the world, and we can do so in a way that vindicates the traditional claims of sense-data theories (sections 5-6). That the basic sense-data language is preferable for talking about our experiences because of its expressive power (section 7). That, with a sense-data language adopted, plausible theories of perception and color can be developed (section 8). And that the sense-data language is the only honest way to face the real skeptical challenge (section 9) and the challenge can actually be met without any cheating (sections 10-11).

The early modern project is not the *only* interesting project in philosophy, nor even the *only* interesting project in epistemology. Yet it *is* interesting, even gripping. History proves that. So does each round of introductory philosophy classes the world over. And the Cartesian skeptical problem, the real skeptical problem, is one of the most disturbing and difficult problems in the entire intellectual landscape.

We can adopt the sense-data language. We can and we should. In fact, in a sense, we *must*. The first-person perspective is forced. We all face the world from the first-person perspective. It is not optional. The early modern epistemological problems are, for this reason, the pure epistemological problems. None of us have direct acquaintance with trees, travelers, and trebuchets. When a philosopher insists that so-called "direct" realism undermines this they have been blinded. Blinded either by philosophical theory or by a fear of skepticism, but blinded all the same. Against these philosophers, the canonical early modern philosophers were right. Not only that, they were *obviously* right. Quantifier variance helps us to clearly understand this.

Sense-data is our only input from the world. It's all that we have, all that we ever will have, and all that we ever could have. The "veil of perception" can be overstated and misunderstood. You might falsely say that the veil is opaque. It is not. It is transparent, but it exists all the same. Our direct acquaintance with sense-data is the epistemological foundation on which we build our overall theory of the world. There is an ongoing, iterated feedback process between data and theory, but *epistemologically* speaking, sense-data are foundational.

This is not to say that sense-data are *metaphysically* foundational. I am no phenomenalist. Phenomenalist theories are inferior to external world theories, explanatorily speaking. Likewise for idealist theories.[132] My point is *epistemological*. I am advocating for *epistemological* empiricism. It is important to realize that my traditional, first-person empiricism differs from what is called "empiricism" in contemporary philosophy. Today's "empiricists" often use a notion of *observable* consequences that is not the first-person, sense-data-based notion.[133] For these philosophers, an observation sentence can be something like "the cat is on the mat".[134] For me, observation sentences are basic experience claims in the sense-data language.

Accepting the sense-data language is part of being a traditional empiricist, but not the whole. Descartes and other rationalists were (in modern terms) sense-data theorists, but they weren't empiricists. This is because they also allowed a substantive role for rational intuitions that sits ill with empiricism. With this, they tried to take armchair reasoning to places that armchair reasoning simply cannot go.[135] For an epistemological empiricist, our sensory states deliver information about reality, but no other mental states do anything similar. There are no rational intuitions through which we gain non-experiential information

---

[132] Though there comes a point at which versions of so-called "objective" idealism only differ verbally from non-idealist theories. For arguments against Berkeley's *theistic* idealism, see Mackie 1982.

[133] See Van Fraassen 1980 for a prominent example. Though the historical and sociological roots of this alternate usage come from the immense and deserved influence of Reid and Quine, respectively.

[134] This divide also corresponds to an internal debate among the logical positivists over "protocol sentences", see Coffa 1991 for discussion.

[135] See section 4 of Warren 2022*b* for criticisms of rationalism.

about any independent world of fact. We have only the sensory data of experience, along with *a priori*, analytic resources with which to build from experience. This sounds meager, but over a century of analytic philosophy has left us with a wide-range of tools.

I call this position analytic empiricism. I prefer this name to "logical positivism" or a "logical empiricism", for two reasons. First, I want to distance myself from strict verificationism, phenomenalism, and the other more questionable doctrines of the Vienna Circle.[136] Second, our analytical tools are not all "logical", so I prefer the broader term. And I prefer "analytic empiricism" to "analytical empiricism", because it is now more common to talk of "analytic philosophy" than "analytical philosophy".

Analytic empiricism is an attempt to complete the traditional empiricist project by using all of the tools and resources of contemporary analytic philosophy that are compatible with empiricism. It is a modern form of traditional British empiricism in the line of Locke, Berkeley, Hume, Mill, Russell, and Ayer. The central plank of analytic empiricism is an adherence to epistemological empiricism as embodied in and limited to the sense-data framework. This paper has been an extended defense of this plank.

From this slender foundation our theory of overall reality must be built. Science, mathematics, ethics, aesthetics, and everything else must fit into this framework or be abandoned. Though long thought dead, traditional empiricism yet lives. In fact, it is in good health. Analytic empiricism represents our last best hope for completing the philosophical project started by the great early modern philosophers. We must build a theory of the world from the data of experience. A theory that can ultimately explain the very existence of sense-data accessing creatures like us. In the end, the snake of philosophy must eat its own tail.[137]

---

[136] The term "logical empiricism" gained prominence via Feigl 1943. Sometimes the phrase is associated with the Berlin branch of logical positivism, see Uebel 2013 for some relevant discussion.

[137] Thanks to Darren Bradley, Justin D'Ambrosio, Matti Eklund, Yu Guo, Eli Hirsch, Uriah Kriegel, Penelope Maddy, and Douglas Stalker. I also want to give special thanks to four dead sense-data theorists: A.J. Ayer, John Mackie, Benson Mates, and Bertrand Russell. I've learned much from all of them, and not just about sense-data. And Ayer and Russell have inspired me since my teenage years. Ayer published 20 books and I have read them all. And though I have read far more than 20 books by Russell, I haven't come close to exhausting his catalog. Thank you both, thank you all; I will try to make you proud.

References

Armstrong, D.M. 1968. *A Materialist Theory of Mind*. London: Routledge and Kegan Paul.

Austin, J.L. 1962. *Sense and Sensibilia*. Oxford: Oxford University Press.

Ayer, A.J. 1940. *The Foundations of Empirical Knowledge*. London: Macmillan.

Ayer, A.J. 1945. "The Terminology of Sense-Data." *Mind* 54(216): 289-312.

Ayer, A.J. 1954. "Can There Be a Private Language?" *Proceedings of the Aristotelian Society, Supplementary* 28: 63-76.

Ayer, A.J. 1956. *The Problem of Knowledge*. London: Macmillan.

Ayer, A.J. 1967. "Has Austin Refuted the Sense-Datum Theory?" *Synthese* 17(2): 117-140.

Ayer, A.J. 1968. *The Origins of Pragmatism*. San Francisco: Freeman, Cooper, & Company.

Ayer, A.J. 1969. *Metaphysics and Common Sense*. London: Macmillan.

Ayer, A.J. 1972. *Probability & Evidence*. New York: Columbia University Press.

Ayer, A.J. 1973. *The Central Questions of Philosophy*. London: Weidenfeld.

Ayer, A.J. 1985. *Wittgenstein*. New York: Random House.

Ayer, A.J. 1992. "Reply to John Fostor." In Lewis Edwin Hahn (ed.), *The Philosophy of A.J. Ayer*. La Salle: Open Court: 198-200.

Beebe, James R. 2009. "The Abductivist Reply to Skepticism." *Philosophy and Phenomenological Research* 79(3): 605-636.

Bennett, Jonathan. 1966. *Kant's Analytic*. Cambridge: Cambridge University Press.

Bennett, Jonathan. 1971. *Locke, Berkeley, Hume: Central Themes*. Oxford: Oxford University Press.

Bennett, Jonathan. 1979. "Analytic Transcendental Arguments." In P. Bieri *et al*. (eds.), *Transcendental Arguments and Science*. Dordrecht: Reidel: 45-64.

Berkeley, George. 2008. *Philosophical Writings* (Desmond M. Clarke, ed.). Cambridge: Cambridge University Press.

Berry, Sharon. 2019. "External World Skepticism, Confidence and Psychologism about the Problem of Priors." *Southern Journal of Philosophy* 57(3): 324-346.

Block, Ned. 1995. "On a Confusion About a Function of Consciousness." *Behavioral and Brain Sciences* 18(2): 227-247.

Bonjour, Laurence. 1999. "Foundationalism and the External World." *Philosophical Perspectives* 13: 229-249.

Bostrom, Nick. 2003. "Are You Living in a Computer Simulation?" *Philosophical Quarterly* 53(211): 243-255.

Broad, C.D. 1923. *Scientific Thought*. London: Routledge & Kegan Paul.

Broad, C.D. 1925. *The Mind and Its Place in Nature*. London: Routledge & Kegan Paul.

Burge, Tyler. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4(1): 73-122.

Burnyeat, Miles F. 1982. "Idealism and Greek Philosophy: What Descartes Saw and Berkeley Missed." *Philosophical Review* 91(1): 3-40.

Cao, Rosa and Jared Warren. forthcoming. "Mental Representation, "Standing-In-For", and Internal Models." *Philosophical Psychology*.

Carnap, Rudolf. 1928. *Der Logische Aufbau der Welt*. Berlin: Weltkreis.

Carnap, Rudolf. 1950. "Empiricism, Semantics, and Ontology." *Revue Internationale de Philosophie* 4(11): 20-40.

Casullo, Albert. 1987. "A Defense of Sense-Data." *Philosophy and Phenomenological Research* 48(1): 45-61.

Chalmers, David J. 1995. "Facing Up to the Problem of Consciousness." *Journal of Consciousness Studies* 2: 200-219.

Chalmers, David J. 2018. "Structuralism as a Response to Skepticism." *Journal of Philosophy* 115(12): 625-660.

Chisholm, Rodrick. 1942. "The Problem of the Speckled Hen." *Mind* 51(204): 368-373.

Chisholm, Rodrick. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge: The MIT Press.

Cornman, James W. 1980. *Skepticism, Justification, and Explanation*. Dordrecht: Reidel.

D'Ambrosio, Justin. 2019. "A New Perceptual Adverbialism." *Journal of Philosophy* 116(8):

D'Ambrosio, Justin. 2022. "An Empirical Solution to the Puzzle of Macbeth's Dagger." *Erkenntnis* 87(3): 1377-1414.

Davidson, Donald. 1973. "Radical Interpretation." *Dialectica* 27(3-4): 313-328.

Dennett, Daniel C. 1991. *Consciousness Explained*. New York: Back Bay Books.

Dinges, Alexander. 2015. "The Many-Relations Problem for Adverbialism." *Analysis* 75(2): 231-237.

Dorr, Cian. 2005. "What we disagree about when we disagree about ontology." In Mark Eli Kalderon (Ed.), *Fictionalism in Metaphysics*. Oxford: Oxford University Press: 234-286.

Dorr, Cian. 2014. "Quantifier variance and the collapse theorems." *The Monist* 97: 503–570.

Dretske, Fred. 1970. "Epistemic Operators." *Journal of Philosophy* 67(24): 1007-1023.

Dummett, Michael. 1956. "Nominalism." *Philosophical Review* 65(4): 491-505.

Dummett, Michael. 1979. "Common Sense and Physics." In Macdonald (ed.), *Perception and Identity*. London: Macmillan.

Eklund, Matti. 2006. "Neo-Fregean Ontology." *Philosophical Perspectives* 20: 95-121.

Eklund, Matti. 2009. "Carnap and Ontological Pluralism." In *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press: 130-156.

Feigl, Herbet. 1943. "Logical Empiricism." in D.D. Runes (ed.), *Twentieth Century Philosophy*. New York: Philosophical Library: 371-416.

Field, Hartry. 1984. "Is Mathematical Knowledge Just Logical Knowledge?" *Philosophical Review* 93(4): 509-552.

Filcheva, Krasimira. forthcoming. "Can There Be a Feature-Placing Language?" *European Journal of Philosophy*.

Fine, Kit. 2005. "Our Knowledge of Mathematical Objects." In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology* 1. Oxford: Oxford University Press: 89-110.

Fumerton, Richard. 1985. *Metaphysical and Epistemological Problems of Perception*. Lincoln: University of Nebraska Press.

Fumerton, Richard. 1992. "Skepticism and Reasoning to the Best Explanation." *Philosophical Issues* 2: 149-169.

Gibbons, Dave and Alan Moore. 1987. *Watchmen*. Number IX, May 1987. DC Comics.

Gifford, Matthew B. 2013. "Skepticism and Elegance: Problems for the Abductivist Reply to Cartesian Skepticism." *Philosophical Studies* 164: 685-704.

Goodman, Nelson. 1946. "A Query on Confirmation." *Journal of Philosophy* 43: 383-385.

Goodman, Nelson. 1951. *The Structure of Appearance*. Cambridge: Harvard University Press.

Goodman, Nelson. 1955. *Fact, Fiction, and Forecast*. Cambridge: Harvard University Press.

Greco, Daniel. 2012. "The Impossibility of Skepticism." *Philosophical Review* 121(3): 317-358.

Grice, H.P. and P.F. Strawson. 1956. "In Defense of a Dogma." *Philosophical Review* 65(2): 141-158.

Hale, Bob. 1987. *Abstract Objects*. Oxford: Blackwell.

Hale, Bob. 2007. "Neo-Fregeanism and Quantifier Variance." *Proceedings of the Aristotelian Society New Series* 107: 375-385

Hale, Bob and Crispin Wright. 2001. *The Reasons Proper Study: Essays Towards a Neo-Fregean Philosophy of Mathematics*. Oxford: Clarendon Press.

Hardin, C.L. 1988. *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett.

Harman, Gilbert H. 1965. "The Inference to the Best Explanation." *Philosophical Review* 74(1): 88-95.

Hawley, Katherine. 2007. "Neo-Fregeanism and Quantifier Variance." *Proceedings of the Aristotelian Society Supplemental* 81: 233-49.

Hawthorne, John. 2006. "Plenitude, Convention, and Ontology." In *Metaphysical Essays*. Oxford: Oxford University Press: 53-69.

Hawthorne, John. 2009. "Superficialism in Ontology." In *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press: 213-230.

Hirsch, Eli. 2002. "Quantifier Variance and Realism." *Philosophical Issues* 12: 51-73.

Hirsch, Eli. 2011. *Quantifier Variance and Realism: Essays in Metaontology*. New York: Oxford University Press.

Hirsch, Eli and Jared Warren. 2019*a*. "Quantifier Variance." In Martin Kusch (ed.), *Routledge Handbook of Philosophy of Relativism*. New York: Routledge: 349-357.

Hirsch, Eli and Jared Warren. 2019*b*. "Quantifier Variance and the Demand for a Semantics." *Philosophy and Phenomenological Research* 98(3): 592-605.

Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.

Huemer, Michael. 2009. "Explanationist Aid for the Theory of Inductive Logic." *British Journal for the Philosophy of Science* 60: 345–75.

Huemer, Michale. 2016. "Serious Theories and Skeptical Theories: Why You Are Probably Not a Brain in a Vat." *Philosophical Studies* 173(4): 1031-1052.

Jackson, Frank. 1975. "On the Adverbial Analysis of Visual Experience." *Metaphilosophy* 6(2): 127-135.

Jackson, Frank. 1977. *Perception: A Representative Theory*. Cambridge: Cambridge University Press.

James, William. 1893. *The Principles of Psychology*. New York: H. Hold and Company.

Johnston, Mark. 1992. "How to Speak of the Colors." *Philosophical Studies* 68(3): 221-263.

Lewis, David K. 1974. "Radical Interpretation." *Synthese* 27(3-4): 331-344.

Lewis, David K. 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74(4): 549-567.

Linksy, Bernanrd. 2014. “Ernst Mally's Anticipation of Encoding.” *Journal for the History of Analytical Philosophy* 2(5): 1-14.

Lipton, Peter. 1991. *Inference to the Best Explanation*. New York: Routledge.

Lipton, Peter. 2001. “Is Explanation a Guide to Inference? A Reply to Wesley C. Salmon.” In G. Hon and S.S. Rakover (eds.), *Explanation: Theoretical Approaches and Applications*. Kluwer: 93-120.

Mackie, J.L. 1976. *Problems from Locke*. Oxford: Oxford University Press.

Mackie, J.L. 1982. *The Miracle of Theism: Arguments for and against the Existence of God*. Oxford: Oxford University Press.

Maddy, Penelope. 2017. *What Do Philosophers Do? Skepticism and the Practice of Philosophy*. Oxford: Oxford University Press.

Martin, M.G.F. 2002. “The Transparency of Experience.” *Mind and Language* 17: 376-425.

Mates, Benson. 1967. “Sense Data.” *Inquiry* 10(1-4): 225-244.

Mates, Benson. 1981. *Skeptical Essays*. Chicago: University of Chicago Press.

McDowell, John. 1982. “Criteria, Defeasibility and Knowledge.” *Proceedings of the British Academy* 68: 455–79.

McGinn, Colin. 1999. *The Mysterious Flame: Conscious Minds in a Material Worl*d. New York: Basic Books.

Moore, G.E. 1909-10. “The Subject-Matter of Psychology.” *Proceedings of the Aristotelian Society* 10: 36-62.

Moore, G.E. 1913-14. “The Status of Sense-Data.” *Proceedings of the Aristotelian Society* 14: 355-380.

Moore, G.E. 1918-19. “The Presidential Address: Some Judgments of Perception.” *Proceedings of the Aristotelian Society* 19: 1–29.

Moore, G.E. 1939. “Proof of an External World.” *Proceedings of the British Academy* 25(5): 273-300.

Okasha, Samir. 2000. “Van Fraassen's Critique of Inference to the Best Explanation.” *Studies in the History and Philosophy of Science* 31: 691-710.

Paul, G.A. 1936. “Is There a Problem About Sense-Data?” *Proceedings of the Aristotelian Society, Supplementary* 15: 61-77.

Peacocke, Christopher. 1983. *Sense and Content: Experience, Thought, and their Relations*. Oxford: Oxford University Press.

Price, H.H. 1932. *Perception*. London: Methuen.

Price, H.H. 1940. *Hume's Theory of the External World*. Oxford: Clarendon Press.

Price, H.H. 1941. “Review of Ayer's *The Foundations of Empirical Knowledge*.” *Mind* 50: 273-293.

Pryor, James. 2000. “The Skeptic and the Dogmatist.” *Noûs* 34(4): 517-549.

Putnam, Hilary. 1971. *Philosophy of Logic*. London: Allen & Unwin.

Putnam, Hilary. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.

Putnam, Hilary. 1983. “Equivalence.” in *Realism and Reason: Philosophical Papers Volume 3*. Cambridge: Cambridge University Press: 26-45.

Putnam, Hilary. 1987. “Truth and Convention: On Davidson's Refutation of Conceptual Relativism.” *Dialectica* 41(1-2): 69-77.

Putnam, Hilary. 2004. *Ethics Without Ontology*. Cambridge: Harvard University Press.

Quine, W.V. 1948. “On What There Is.” *Review of Metaphysics* 2(5): 21-38.

Quine, W.V. 1951. “Two Dogmas of Empiricism.” *Philosophical Review* 60(1): 20-43.

Quine. W.V. 1957-8. “Speaking of Objects.” *Proceedings and Addresseses of the American Philosophical Association* 31: 5-22.

Quine, W.V. 1960*a*. *Word and Object*. Cambridge: The MIT Press.

Quine, W.V. 1960*b*. “Variables Explained Away.” *Proceedings of the American Philosophical Society* 104(3): 343-347.

Reid, Thomas. 1983. *Inquiry and Essays* (Beanblossom and Lehrer, eds). Indianapolis: Hackett.

Rinard, Susanna. 2017. "Skepticism and Inference to the Best Explanation." In McCain and Poston (eds.), *Best Explanations*. Oxford: Oxford University Press: 203-215.

Robinson, Howard. 1994. *Perception*. New York: Routledge.

Roche, William and Elliot Sober. 2013. "Explanation is Evidentially Irrelevant; or, Inference to the Best Explanation meets Bayesian Confirmation Theory." *Analysis* 73: 659-668.

Russell, Bertrand. 1912. *The Problems of Philosophy*. London: Williams and Norgate.

Russell, Bertrand. 1914*a*. "The Relation of Sense-Data to Physics." *Scientia* 16: 1-27.

Russell, Bertrand. 1914*b*. *Our Knowledge of the External World*. Chicago: Open Court.

Russell, Bertrand. 1927. *The Analysis of Matter*. New York: Harcourt, Brace.

Russell, Bertrand. 1959. *My Philosophical Development*. New York: Simon and Schuster.

Ryle, Gilbert. 1954. *Dilemmas*. Cambridge: Cambridge University Press.

Schiffer, Stephen. 1996. "Language-created Language-independent Entities." *Philosophical Topics* 24(1): 149-167.

Schiffer, Stephen. 2003. *The Things We Mean*. Oxford: Oxford University Press.

Searle, John. 2015. *Seeing Things As They Are: A Theory of Perception*. Oxford: Oxford University Press.

Sellars, Wilfrid. 1975. "The Adverbial Theory of the Objects of Sensation." *Metaphilosophy* 6: 144-160.

Sellars, Wilfrid. 1997. *Empiricism and the Philosophy of Mind*. Cambridge: Harvard University Press.

Sider, Theodore. 2009. "Ontological Realism." In *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press: 384-423.

Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.

Slote, Michael A. 1970. *Reason and Scepticism*. London: Allen & Unwin.

Stalker, Douglas (ed.). 1994. *Grue! The New Riddle of Induction*. Chicago and La Salle: Open Court.

Strawson, P.F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.

Taylor, Charles. 1979. "Sense Data Revisited." In G.F. Macdonald (ed.), *Perception and Identity: Essays Presented to A.J. Ayer with His Replies*. Ithaca: Cornell University Press: 99-112.

Thomasson, Amie L. 2009. "Answerable and Unanswerable Questions." In David J. Chalmers, David Manley, and Ryan Wasserman (eds.), *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press: 444-471.

Turner, Jason. 2011. "Ontological Nihilism." In Karen Bennett (ed.), *Oxford Studies in Metaphysics* 6: 3-54.

Tye, Michael. 1984. "The Adverbial Approach to Visual Experience." *Philosophical Review* 93(2): 195-225.

Uebel, Thomas. 2013. ""Logical Positivism" — "Logical Empiricism": What's in a Name?" *Perspectives on Science* 21(1): 58-99.

Van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Clarendon Press.

Van Fraassen, Bas C. 1989. *Laws and Symmetry*. Oxford: Clarendon Press.

Van Inwagen, Peter. 1990. *Material Beings*. Ithaca: Cornell University Press.

Vogel, Jonathan. 1990. "Cartesian Skepticism and Inference to the Best Explanation." *Journal of Philosophy* 87(11): 658-666.

Warren, Jared. 2015. "Quantifier Variance and the Collapse Argument." *Philosophical Quarterly* 65(259): 241-253.

Warren, Jared. 2016*a*. "Sider on the Epistemology of Structure." *Philosophical Studies* 173(9): 2417-2435.

Warren, Jared. 2016*b*. "Internal and External Questions Revisited." *Journal of Philosophy* 113(4): 177-209.

Warren, Jared. 2021*a*. "This Quintessence of Dust – *Consciousness Explained*, at Thirty." *Philosophical Papers* 50(1-2): 281-308.

Warren, Jared. 2021*b*. "Defending Understanding-Assent Links." *Synthese* 199(3-4): 9219-9236.

Warren, Jared. 2022*a*. “Quantifier Variance, Semantic Collapse, and “Genuine” Quantifiers.” *Philosophical Studies* 179(3): 745-757.

Warren, Jared. 2022*b*. *The* A Priori *Without Magic*. Cambridge: Cambridge University Press.

Weisberg, Jonathan. 2009. “Locating IBE in the Bayesian Framework.” *Synthese* 167: 125-143.

Williams, Donald C. 1947. *The Ground of Induction*. New York: Russell and Russell.

Williams, Michael. 2010. “Descartes' Transformation of the Sceptical Tradition.” In Richard Bett (ed.), *The Cambridge Companion to Ancient Scepticism*. Cambridge: Cambridge University Press: 288-313.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

Williamson, Timothy. 2006. “Conceptual Truth.” *Aristotelian Society Supplementary* 80(1): 1-41.

Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. New York: Wiley-Blackwell.

Wright, Crispin. 1983. *Frege's Conception of Numbers as Objects*. Aberdeen: Aberdeen University Press.