UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Having Things in Mind

Permalink
https://escholarship.org/uc/item/96g7j8xv

ISBN
9798173788061

Author
Levac, Madeleine

Publication Date
2026-06-10

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Having Things in Mind

By

Madeleine Levac

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Philosophy

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor John G. MacFarlane, Co-Chair

Professor Michael G.F. Martin, Co-Chair

Professor Line Mikkelsen

Summer 2026

Having Things in Mind

Copyright © 2026

By

Madeleine Levac

Abstract

Having Things in Mind

by

Madeleine Levac

Doctor of Philosophy in Philosophy

University of California, Berkeley

Professor John G. MacFarlane, Co-Chair

Professor Michael G.F. Martin, Co-Chair

Seeing a persimmon on a nearby branch, you begin to have thoughts about it. You pick the persimmon and put it in your pocket. In this sequence of events, you are related to a particular fruit in various ways. But whereas the presence of a persimmon is clearly indispensable to the action of *picking a persimmon* or the state of *having one in your pocket*, philosophical reflection draws us away from according objects the same status relative to perception and cognition.

Imagine a variant on this case in which you merely hallucinate a persimmon on the branch above you. Reaching out, your hand grasps empty air. You cannot truly claim to have *picked* a persimmon, or to be *standing underneath* one. But can you say that you *saw* one? That you *wanted* to pick it? Here we might answer "yes," and with increasing conviction: these questions reflect a concern with the subjective perspective, which the envisaged change in your external circumstances need not touch.

It is tempting to retreat, on this basis, to an internalist understanding of the relevant mental states. I aim to inoculate the reader against that temptation. Following Evans and McDowell, I argue that some of our thoughts about things in the world are fundamentally relations to—not merely representations of—the things that they concern. I do so by developing an underexplored idea of Evans's, namely that some of our ordinary mental state ascriptions occur in the bounds of a pretence. When this suggestion is appropriately developed and applied, the result is a radical improvement in both the relative standing and the independent appeal of a conception of singular thought as object-dependent.

Table of Contents

Acknowledgements....ii
Introduction....iii
Chapter 1: Prior's Challenge....1
Chapter 2: The Task Ahead....12
Chapter 3: Minds and Make-Believe 30
Chapter 4: Applications....,....48
Chapter 5: From Thought to Perception....66
Chapter 6: Meeting the Challenge....98
References....113

Acknowledgements

In the process of writing my dissertation I have benefitted from the talents and the energies of many people. I first thank my advisors, John MacFarlane and Michael Martin. John has weathered years of chatter about *The Varieties of Reference*, beginning when he indulged my desire to do a qual with him entirely about that book. John is a wonderful interlocutor, and his scepticism about some of the claims in this dissertation—always tactfully expressed and thoroughly justified—has resulted in a much clearer and less dogmatic text than I might otherwise have produced. From Mike I have learned more about philosophy than I could reasonably have hoped. He discerns in an instant what a mere mortal like myself can take weeks to grasp; efforts to catch up to him are always well rewarded, even if never quite successful. Mike's uncommon degree of insight also manifests in his advising, where it is complemented by an exceptional devotion to his students. His sensitivity to my interests and able shepherding of my project have only become fully evident to me in retrospect, and his assurances that things would come together eventually (pronounced with what seemed to me to be great conviction) were instrumental in securing that outcome.

I also want to extend special thanks to Kendall Walton. Ken happened to join us at Berkeley right around the time that he was becoming a key figure in my dissertation. His evident kindness and sincere interest in the thoughts of graduate students helped to revive my desire to engage philosophically at a moment when I was feeling rather beaten down. I am also grateful for the ways that conversations with Ken improved my understanding (though still imperfect, I'm sure) of his ideas.

Among the faculty at Berkeley, I thank John Campbell, Geoffrey Lee, and Seth Yalcin for exchanges that shaped the content of this dissertation, as well as Olivia Bailey and Andreja Novakovic for other forms of positive influence. Many thanks to the office staff that ensured the smooth running of the department during my time there, especially Janet Groome, whose recent departure is a major loss.

I have been sustained personally and intellectually by friends, fellow graduate students (current or former), and a significant number of people who meet both descriptions. Thank you in particular to Greyson Abid, Esther Adelstein, Katherine Ammirati, Austin Andrews, Sophia Arbeiter, Mike Arsenault, Monika Chao, Luna Cheng, Katie Coyne, Jordan Dopkins, Micah Dubreil, Jennifer Eaman, Caroline Erickson, T. Haddow, Adam Kochanski, Jenny Marsh, Emily Perry, Amy Rinaldi, Molly Rose-Williams, Patrick Ryan, Pia Schneider, Umrao Sethi, Joy Shim, Gabriel Siegel, Warren Thimothe, Sarah Vernallis, and Jamie Yin. I am grateful to Emma Lasky for, among other things, the gift of noise-cancelling headphones.

Thank you to Lauren Massell for her outsized contributions to my psychological and emotional well-being, and to David and Elizabeth Kaleugher for hosting me on numerous escapes to Colorado. I thank Michael and Nancy Kaleugher for their encouragement and their myriad forms of support. J'exprime mon appréciation perpétuelle pour mon oncle Gérald Levac, qui me manque beaucoup. Most of all, I thank my mother, Pamela Levac, for the person she is, the haven she provides, and for the autonomy she has granted me from an early age over my life and education.

Introduction

Some of our activities require sustained interaction with particulars in the world around us. In throwing a vase, for instance, one engages repeatedly with a piece of clay. Remove the clay, and the potter can at best imitate her usual actions. There is no genuine throwing of a vase in the absence of suitable materials, and little sense can be made of a person who persists in the attempt though she knows those basic external requirements aren't met.

Philosophical orthodoxy has it that the activities of the mind are different in this regard: while we can think about the very particulars that we act on with our bodies, thinking about a thing is not dependent on its object in anything like the way that throwing a vase is dependent on the clay. To conceive of a mental state in this way is to hold, in contrast to the case of embodied action, that it could occur with or without whatever it purports to be directed upon—that it could occur in either case and yet still be a state of the same fundamental kind.

This conception is my primary critical target in this dissertation. I claim that some of our thoughts are fundamentally relations to the mind-independent things that they concern. Establishing that there are any such cases is sufficient to show that the mind and the non-mental world are interpenetrating. If that's right, our minds cannot be indifferently transplanted from one environment to another in the way that philosophers since Descartes have often supposed: you would not be psychologically the same if you were a brain in vat or if you were living in a qualitative duplicate of our actual world, for instance.

This metaphysical claim has important epistemological and experiential consequences, some of which will be objects of sustained attention in the pages that follow. It also vindicates our common-sense, pre-philosophical understanding of ourselves and of our situation. In this respect, the view I defend here breaks with mainstream analytic philosophy of mind. Many of us are inured to the possibility that our most foundational beliefs are wildly incorrect, the products of naïve optimism or intellectual laziness. We might think that it is in the essence of philosophical inquiry that it be alienating or revisionary. But the philosopher and the ordinary person need not be at odds in order for the activities of the former to be worthwhile. Merleau-Ponty describes a philosophy "for which the world is always 'already there' prior to reflection, [...] and whose whole effort is to restore this naïve contact with the world in order to give it, at last, a philosophical standing." (1945, *i*; my translation) What he intends as a characterization of the phenomenological tradition might be applied more widely. Philosophy can surprise and unsettle, but it can also deepen our understanding of what we already, at some level, know. I hope that this text can serve as a reminder of the latter facet of our enterprise.

Overview of chapters

Chapter 1 introduces the initiating problem of the dissertation. What is it to think about a thing—a particular, mind-independent constituent of reality? A challenge issued by Prior (1971) clarifies both the question and the main ways of answering it:

> *a)* X's thinking of Y constitutes a relation between X and Y when Y exists, but *(b)* not when Y doesn't; but *(c)* X's thinking of Y is the same sort of thing whether Y exists or not. Something plainly has to be given up here; what will it be? (130)

I consider and reject several attempts to show that the problem represented in these lines is illusory. I then outline the responses to the challenge that we arrive at by giving up *(a)*, *(b)*, and *(c),* respectively. The *representational strategy* gives up the idea that thinking of a thing is fundamentally a relation to that thing. The *ontological strategy* posits special kinds of entity or invokes different grades of being, with the effect that thinking of Y turns out to be a relation whether Y (prima facie) exists or not. The *disjunctive strategy* accepts that there are two fundamentally different mental states here, insisting on the correctness of *(a)* and *(b)*.

Chapter 2 lays the ground for a resolution to Prior's challenge by way of the disjunctive strategy. I first present a more cautious restatement of the claims that Prior assembles, which helps us see that the prospects for the disjunctive strategy are better than they appear when we take the challenge at face value. Next, I situate the disjunctive strategy in its philosophical context. I explain the forces that have pushed it towards the margins and discuss the family of views in the philosophy of mind, epistemology, and philosophy of action to which it belongs. I then indicate the general structure of the implementation of the disjunctive strategy that I pursue in subsequent chapters.

Chapter 3 in a way constitutes the core of the dissertation, so I summarize it here in particular detail. It introduces an account of singular thought as object-dependent, as well as a series of objections that are commonly levelled against that view. Those objections concern its handling of what we can call *empty cases:* cases where it is for the subject as if she has a singular thought, though there is no object that she thinks about. The advocate for the object-dependent conception diagnoses the appearance of singular thought, understanding, and communication in such cases as *merely* apparent, denying that there is a thought that the subject entertains, a proposition she expresses in giving voice to her subjective state, a content for the hearer to understand—claims that may strike us as wildly implausible.

The primary purpose of the chapter is to excavate from Evans's (1982) *The Varieties of Reference* a treatment of the empty case whose precise character, range of applications, and explanatory power have not been fully appreciated. The key insight, inspired by a pair of papers by Walton (1973, 1977) that predate Evans's lamentably early death, is that some apparently ordinary mental state ascriptions are in fact produced and taken up internal to a pretence. Pretence, for Walton, is imagining scaffolded by props—features of the real world that combine with rules explicitly stipulated or implicitly assumed by the participants to determine what is "make-believedly" the case, i.e., true-in-the-pretence.

The employment of this idea that will primarily concern us here has the following key features. It is a form of *existentially creative* pretence, in which one pretends that things exist which in fact do not. Its central props are *psychological* rather than concrete: mental representations of a kind that Evans calls "informational states" (1982, 122). Informational states embody information that at least purports to derive, and typically does derive, from a single object; they are individuated by their sources, not merely by their representational content. They contrast with singular thought, per Evans, in being both object- and belief-independent. They may be the product of misinformation as well

as genuine information (here the term "information" will be used as neutral between the two).

Suppose we have an informational state of the *mis*informational variety, which is not genuinely "of" any individual (compare the sense in which a photograph is "of" the sitter), but which nonetheless represents one. If we pretend that the informational state is veridical, we *by that very fact* make it the case that an object is introduced in the pretence—introduced as that from which the information in question derives. This object can be redeployed, within the pretence, as a basis for singular reference, singular thought, and the communication of singular propositions. How does this bear on the problems that the empty case introduced? The suggestion is that a failure to recognize this pretending for what it is explains our robust impression of reference, thought, communication, and understanding even in the absence of a (real) object of thought.

In chapter 4, I make use of this proposal in two ways. First, in connection with disagreements about the necessary and sufficient conditions on tokening a singular thought, and particularly about the role of acquaintance. Second, in connection with the phenomenon known as "intentional identity," in which multiple attitudes are judged to be *about the same thing*—on non-descriptive grounds—even though their putative object does not exist. There are aspects of the literatures on both topics that stand to be reconsidered in light of the novel account of the empty case that was put on the table in chapter 3. These applications provide an opportunity both for the further elucidation of the claims of that chapter, and for a clarification of the relationship between the representational, ontological, and disjunctive strategies introduced in chapter 1.

Chapter 5 is an attempt to settle a long-running debate about the meaning of the verbs of perception with the aid of the general strategy developed in chapter 3. Perceptual reports (e.g., "She hears the whistle," "They see Frankie") are widely acknowledged to have two uses, only one of which commits the speaker to the existence of the object said to be perceived. The existing literature offers us two main ways of accounting for this fact, neither of which has managed to emerge victorious. I trace this stalemate to the neglect of a third interpretation of what we are doing when we (for instance) say, knowing full well that there is no dagger suspended before Macbeth, that Macbeth sees a dagger.

On my account, which is an analogue for sensory experience of the proposal I developed in the context of cognition in chapters 3 and 4, such utterances are apt despite failing to express a proposition at all. The distinctive way that particulars figure in perceptual phenomenology—a phenomenology inherited by hallucinations, since the subject may be unable to distinguish these from perceptions on a purely introspective basis—demands a certain kind of description from us. As that description involves singular reference to the object of the perceptual relation, it's one that, in a case of hallucination, we lack the materials to make. Hence the motive to pretend that hallucinating subject is perceiving and make an ordinary ascription internal to the resulting pretence. If that's the right account of what we are doing, then we should not be surprised that the debate about the perceptual verbs has enjoyed so generous a lifespan: the various parties have differed on the question of what is asserted and how, but the assumption that these pronouncements are genuine assertions has not been challenged.

In chapter 6, I consider explicitly some themes that have emerged only in passing up to that point. I also make my final case for adopting the disjunctive strategy as a

response to Prior's challenge. As I propose to understand the disjunctive strategy, it does not require us to "give something up" after all: the initial impression that it is on a par with the representational and ontological strategies in this regard is dispelled once we clarify the status of Prior's third claim (that X's thinking of Y is the same sort of thing whether Y exists or not) and uncover the explanatory dividends of an appeal to pretence.

The framework that I adopt from Evans has the striking consequence that some of the realities we acknowledge—subjective realities, specifically—are not amenable to direct description. That's one way into a broader point about the subjective point of view that pervades this dissertation. Much of our insight into the thoughts and experiences of others issues from our grasp of the objective, mind-independent realities that they think about and occupy a perspective on. Achieving insight into someone's conscious states, on this picture, is not a matter of penetrating to a private realm of mental representations; it is a matter of considering the world from her position.

Chapter 1: Prior's Challenge

Some of our thoughts, intentions, and feelings are directed towards things in the world around us. The very entities that we encounter in perception or act on with our bodies can be objects of contemplation, desire, or fear. These commonplaces retain their hold on us in all but the most sceptical of moods: they report some of the facts about the mind that, as philosophers, we feel called to explain. In this first chapter, I draw attention to a fundamental choice point in theorizing about the relationship between thought and its objects.

1.1 Background

The mental states just alluded to are distinguished by their *intentionality*: they are about, or directed upon, something. (This definition is deliberately vague, for reasons that will soon become clear.) Franz Brentano introduces this term to the contemporary philosophical lexicon in the following well-known passage:

> Every mental phenomenon is characterized by what the Scholastics of the Middle Ages called the intentional (or mental) inexistence of an object, and what we might call, though not wholly unambiguously, reference to a content, direction toward an object (which is not to be understood here as meaning a thing), or immanent objectivity. ([1874] 1995, 88)

This statement is historically important, but also, it must be noted, idiosyncratic. Few would deny that *some* mental states are intentional in the relevant sense, but the suggestion that intentionality is the mark of the mental is controversial.[1] In addition, Brentano's definition shows the influence of early views that he himself later rejected. He held that the targets of not only psychological but of all scientific investigation are mind-dependent phenomena, causal products of a more fundamental reality that is inaccessible to us.[2] This is what induces him to say that the object of intentional state is "not a thing" and to speak of intentional *in*existence (by which he means that the object of a thought is internal to the thought itself).[3]

If instead we take it that questions about the objects of attitudes are at least sometimes rightly answered by referring or pointing to things that exist independently of the mind, we have to wrestle with an issue that doesn't arise on Brentano's original picture. After all, questions about the objects of attitudes are *also* sometimes rightly answered with an expression that doesn't single anything out.[4] A credulous child wants to adopt Pegasus as a pet—there's no such creature. Ponce de Leon seeks the Fountain of

---

[1] The phenomenal is another feature regularly cited in this connection, often alongside rather than to the exclusion of the intentional. See (Rorty 1979, chap. 1 sec. 2, esp. 22–24) for an efficient, albeit highly sceptical, introduction to this two-criterion picture.

[2] One expression of this view is in his discussion of the relationship between psychology and the natural sciences in book 1, chapter 1 of *Psychology as an Empirical Standpoint.* See esp. (Brentano [1874] 1995, 14).

[3] Crane (2006) discusses these points at much more length.

[4] In setting things up this way I take inspiration from (Anscombe 1965).

Youth—it doesn't exist. You fume about your ex's cool new partner—a fiction concocted to make you jealous.[5] So we have our own reasons to wonder about the significance of the term "object" as it appears in such phrases as "object of thought," and hence also to wonder what the "directedness" or "aboutness" of intentional mental states consists in.

1.2 Three appealing claims

Intentional states are characterized by their reference to an object. That object need not exist. Despite the air of paradox, these two claims are consistent if instituted in the right way. One might for instance hold, as A.N. Prior (1971) points out:

> That thinking of an unreal object is quite a different sort of activity from thinking of a real one.
>
> That in thinking of anything at all, we thereby put ourselves into a relation, not with *that thing*, but only with an 'idea' or what-have-you which in favourable cases may 'represent' a real thing but in unfavourable ones does not.
>
> That there are strong and weak kinds of reality or being, such that all objects of thought whatever possess at least the weak sort, while only a favoured sub-class possess the strong sort. (127)

The problem lies, not in finding a way to be consistent, but in deciding *which* consistent thing to say, and then saying it convincingly. Yet some philosophers shirk the burden of decision altogether, or so Prior complains in the essay where these examples appear.[6] That essay of Prior's terminates in the following lines:

> *a)* X's thinking of Y constitutes a relation between X and Y when Y exists, but *(b)* not when Y doesn't; but *(c)* X's thinking of Y is the same sort of thing whether Y exists or not. Something plainly has to be given up here; what will it be? (1971, 130)

Call this *Prior's challenge.* It will serve a dual purpose here. First, it constitutes a crisp statement of the problem that I am setting out to solve. Second, it imposes some initial structure on the space of possible responses to that problem. In the remainder of this chapter, I first establish that the challenge is a genuine one, and then introduce in more detail the three possible lines of response to it.

1.3 Making the challenge precise

I take it that the claims Prior lays out for us have considerable intuitive appeal. As stated, however, they leave open various ways that the challenge might be sidestepped. A

---

[5] The problem here is not the possibility of thought about Pegasus *qua* creature of myth, or the new partner *qua* carefully crafted fiction. These are plausibly abstract artefacts. The problem is the possibility of thought about Pegasus *qua* winged horse or the new partner *qua* actual person, whether that occurs in the course of an imaginative flight of fancy or because of sincere confusion about what reality contains.

[6] His particular targets include (Reid 1785), (Findlay 1961), and (Anscombe 1965).

first issue is that the triad is not obviously inconsistent. This is a result of the imprecise language in *(c)*: what it is for two phenomena to be "the same sort of thing," and how significant a relationship between them this imposes, depends entirely on the "sort" in question.

It's clear enough for present purposes what Prior has in mind. The claim is that there is *one* thing we do, "thinking of," though we do it under (extrinsically) different circumstances; that the mental states described in *(a)* and *(b)* are states of the same *fundamental* kind.[7] What I want to focus on now are three more substantial responses to the challenge—responses that each, in one way or another, have the potential to defuse it altogether. The first pushes back on the expectation that intentionality would be relational (in other words, that *(a)* is intuitively true). The second questions the implicit assumption that there is an interesting relationship between the cases at issue in *(a)* and *(b)*. The third traces the appearance of a problem to confusion about the way that we talk about the intentionality of mental states.

1.3.1. Pressure on the relational intuition

Let's restrict our attention momentarily to those cases where the object of a subject's thought exists. Why suppose that thinking of an object is a way of being related to it, on the model of, say, holding an object in one's hand? One candidate reason is that the structure of our ordinary attitude reports shows that we conceive of it in this way.

"X thinks of Y," and other objectual attitude reports ("X fears Y," etc.), unquestionably have a relational syntax: two names or noun phrases are linked by a verb. This invites a correspondingly relational account of their semantics. Namely: each of the expressions substituted for "X" and "Y" have the role of referring to some entity (or group of entities), and the verb describes a way that the first is involved with the second. The circumstances under which we deem these sentences true, however, cast doubt on this proposal. Like "searches for" or "paints," many psychological verbs are such that simple sentences constructed from them strike us as true irrespective of whether the expression in object place refers.

To see that the underlying reality might be more complicated, consider the following. One might similarly assume that a claim like

(1) The dog chased a bug.

has the logical form of a binary relation.[8] According to a Davidsonian treatment of action sentences, however, we should analyze (1) in terms of an event adverbially modified: for some event *e*, *e* was a chasing, and the dog was the agent of *e*, and a bug was (as the hapless recipient of the action) the "theme" of *e*.[9] Intensional transitive verbs, attitude verbs among them, might well call for the same treatment.[10] But a semantics along these lines puts events, not relations between the subject and the object of thought, at center

---

[7] In the end a lot does turn on this clarification; see §2.1.2.1.
[8] As suggested for instance in (Montague 1973), (Thomason and Stalnaker 1973).
[9] See (Davidson 1967) for the original proposal, and (Parsons 1990) for the refinement of the view via the introduction of semantic θ-roles that I draw on here.
[10] (Forbes 2006, chap. 5, sec. 4).

stage; here the noun phrase that follows the verb is given the role of characterizing a theme, rather than that of identifying an object that the subject is related to. It's a further question whether the *existence* of the theme is essential to the event—whereas it's not a further question whether the existence of the second term is essential to a binary relation.

The possibility that attitude reports might be correctly interpreted in this way means that we cannot confidently rest the weight of Prior's challenge on the superficial form of such phrases as "thinking of" or "is about" (we might add that the founding of substantive philosophical claims on cursory linguistic observations is ill-advised in general). If *(a)* has a genuine hold on us, it must be for deeper reasons. Here, then, is a first respect in which the challenge needs buttressing: why think intentionality is a relation in any interesting sense?

1.3.2 On the connection between (a) and (b)

Let's set the concerns of the last section to the side for a moment in order to raise another worry. Prior's challenge brings together a pair of psychological situations, in *(a)* and *(b)*, in a way that implies that these are of special relevance to each other. In the first, the subject (X) thinks about an individual (Y). That this is *prima facie* a relation suggests that what's at issue is a so-called "singular thought," a kind of thought whose object is (or objects are) determined non-descriptively.[11] In the second situation, there is a subject (X again) and a merely putative entity (Y-that-doesn't-exist) which nonetheless is cited as the object of X's thought.

We are supposed to be pulled one way and then the other about the nature of intentional states by these two cases—we are supposed at least to be strongly tempted by the prospect of a unified account of both. But one might ask whether it makes sense to pit them against each other in the first place, considering that our ascriptions of intentional mental states take various forms and that attitude verbs can be followed by a quite heterogeneous range of expressions. For example:

> John wished for a sloop [non-specific reading].[12]
>
> She's thinking about bicycles [in general], and what a clever invention they are.
>
> Sarah is angry at the person responsible [whomever it is / if there is one].

Focusing on examples like these, we are unlikely to be drawn towards a claim like *(a)*. Yet it's not obvious that these ascriptions are any less paradigmatic than those that implicate a particular thing—at the very least, establishing as much would require some argument. Why not understand an ascription like "Rieva loves Pegasus" on the model of "Rieva loves horses," for instance?[13] Perhaps the various attitudes at issue are the same

---

[11] I introduce this concept at greater length in chapter 3.
[12] (Quine 1956).
[13] Sainsbury (2010) raises a version of this objection that combines it with the issue in my previous subsection, §1.3.1.

sort of thing simply in that they share in the property of intentionality; some happen to be relational in character, some aren't, and there is no problem with that.

The salience of this objection reflects a disanalogy between the issue of empty representations as it appears here and in the philosophy of perception. On Prior's framing, the two cases are parallel:

> Perception: Perceiving a persimmon is a relation between you and that persimmon. A hallucination as of a persimmon cannot be a relation to a persimmon, since by hypothesis no persimmon is playing a causal role. Are perceiving and hallucinating the same sort of thing, or not?
>
> Intentionality: Thinking of a persimmon you just ate is a relation between you and that persimmon. A thought about a mythical persimmon cannot be a relation to a persimmon, since there is no persimmon to be the object of the thought. Are thinking of a real and of a non-existent object the same sort of thing, or not?

This obscures an important difference in dialectical landscape across the two cases.

The debate about the metaphysics of perception is shaped by the causal argument from hallucination. That argument trades on a plausible consequence of the subjective indiscriminability of hallucinations and veridical perceptions—a consequence, in other words, of the fact that these may not be distinguishable just by introspection.[14]Suppose that perceiving is indeed a relation to things in the environment. We can conceive of a hallucination that a subject cannot tell apart (drawing on a single perceptual modality, at a time) from a token veridical perception. It's plausible that this indiscriminability means that the phenomenal facts are the same across both cases. It's also plausible that phenomenal facts are locally supervenient, determined by what is going on with the brain. But then a state of the subject is causally sufficient for the experience she is having when she perceives. If we deny that there is action at a distance (i.e., causal influence without contact or a medium of transmission), the fact that the perceptual and the hallucinatory states differ in their distal causes cannot distinguish them. These materials form the basis of a challenge to the idea that perception is *fundamentally* a relation to things in the environment.[15]

We may likewise suppose that cases in which one thinks about something, and cases in which there is nothing about which one thinks, can be impossible to tell apart from the thinker's perspective. But what inclines us to say this? Probably not a conviction that the two mental states have identical phenomenology.[16] This is problematic because it

---

[14] A *veridical* perception is a case of the subject's being aware, through their senses, of the world around them, which appears as it in fact is. The contrast is with experiences like hallucination and illusion. Arguably ordinary usage of the term "perception" is such that it is veridical by definition; the addition of "veridical" is a precaution against ambiguity.

[15] See (Martin 2004) for a more detailed presentation of these issues.

[16] Indeed, it's doubtful whether thoughts are rightly described as having a phenomenology at all. Much of the disagreement in the literature on this topic concerns whether those who claim that there is cognitive phenomenology are just picking up on accompanying sensory phenomenology. Against this diagnosis, see for instance (Siewert 1998, chap. 8), (Pitt 2004). For some representative expressions of scepticism about cognitive phenomenology, see (Nelkin 1989), (Tye 1995: 4), and (Carruthers 2005).

deprives us, in the case of intentionality, of one prominent kind of reason we do have to group perception and hallucination together; namely, a presumed overlap at the "mind" end of their causal trajectories. Hence our second problem. If we really ought to expect that the cases in *(a)* and *(b)* pose a problem for each other, as the causal argument suggests that perception and hallucination do, it remains to be said why.

1.3.3 Grammatical observations

One of the intended targets of Prior's challenge is G.E.M. Anscombe's (1965) paper on intentionality in perception. The way that he presents her view suggests that he has not fully absorbed its lessons, however, for it offers us another way that his problematic trio of claims might be circumvented.

Anscombe wants to dispel the assumption that when we speak of the intentional object of an attitude, we must be speaking of a kind of thing. She maintains instead that the concept of an intentional object is a "grammatical notion" (166): a concept, like that of a direct object, that serves to elucidate the structure of our language.[17] In our ascriptions of intentional states, we identify an attitude by way of a psychological verb and follow that verb with a noun phrase. That noun phrase may or may not refer, but it clearly plays a role in characterizing the subject's mental life, just as the verb does. That our ascriptions have this structure, and that they have it for some *reason*, is what a statement like "every intentional state has an object" ought to be understood as expressing.

Anscombe gives us some insight into that reason with the contrast she draws between intentional and "material" objects, illustrated by means of a case in which "[a] man aims at a stag [,] but the thing he took for a stag was his father, and he shoots his father" (167). By citing the intentional object of the aiming, we give a "description under which" its agent would recognize his own action (166): saying that he aimed at a stag, we faithfully represent what he thought he was doing. We may, however, instead wish to emphasize that what he *in fact* aimed at—the thing in his sights—was his father. In this case we identify a material object of the aiming: we say what thing, if any, an attitude or intention is directed towards.

These labels aren't the most perspicuous, but the suggestion comes to this. There is a difference in the truth-conditions of intentional attitude ascriptions according to what we intend to express in making them. Some of those ascriptions are intended to describe how a subject is in fact linked up with other things in the world. Others are intended to describe a subject's perspective from within. Ascriptions of the first kind are true if they accurately describe the total situation; ascriptions of the second kind are true if faithful to the subject's experience of it.

Against this background, the fact that we say "X is thinking of Y" in cases where Y exists and in cases where it doesn't is no evidence of an underlying inconsistency in how we understand intentional states. The question whether thinking of Y is or isn't a relation to a thing just looks misplaced. We'll want to say more about what is being communicated in each case, and how, and about the relationship between the two uses. But those are linguistic and conceptual issues; the metaphysical question that Prior's

[17] More on this claim in §2.1.1 below.

challenge appears also to present us with has, or so the objection runs, been explained away.

1.4 Responses

I'll address these ways of defusing Prior's challenge in reverse order. First, Anscombe's observations do not rule out that there is a genuine problem about the nature of intentional states, even about our ordinary concepts of those states. Allow for the sake of argument that she succeeds in showing that our ordinary talk of attitudes and their objects does not commit us to inconsistency. Suppose that she does explain away the seeming mystery as to what intentional objects are. Prior's challenge is, first of all, not a challenge about what intentional objects could possibly be. More crucially: one can still ask *what state of affairs* we are describing when we make the two uses of intentional attitude ascriptions that Anscombe distinguishes. What is it to think about an object, when that object is among the things that exist? What situation obtains when a subject is aptly described as thinking about *Y*, though "*Y*" is not a name for anything? What is the relationship between these two conditions? These are questions that grammatical observations simply do not resolve.

The second objection to the challenge concerned the assumption, in Prior's setup, that *X's thinking of Y when it doesn't exist* belongs alongside *X's thinking of Y when it does exist*—and is not simply to be treated as one of several kinds of intentional ascriptions that don't make reference to any specific thing. We saw that subjective indiscriminability is not as natural a way of binding this pair of cases together as it is in the context of sensory experience. I propose that we attend to *particularity* instead.[18] It's by reflecting on the difference that particulars make to our thought that we discover what the two kinds of case presented in the challenge have in common, and why it makes sense to consider them together. In a way this is a point that it will take me the rest of the dissertation to make; for the purposes of this stage-setting chapter, I'll just say enough to get the problem off the ground.

In addition to thinking about *an F* or *Fs* or *the F* or thinking that there is some object such that… (etc.), there is thinking about an individual *as* such, that is, not merely as the satisfier of some descriptive condition.[19] There are some lines of reasoning, some attitudes, behaviours, and reactions, that cannot be made sense of in terms of a subject's entertaining a description—not even a description that happens to have a unique satisfier. At time *t1* a child has a hamster that is fluffy and beige. At time *t2* the child has a hamster that is fluffy and beige. Why is she upset? Because in the meantime her original pet hamster died and was covertly replaced with a qualitative duplicate. To register this occurrence, we need to pick out the original hamster relationally. The reason for the child's sadness or disappointment cannot be articulated in qualitative terms, because qualitatively (we can imagine) nothing has changed. Likewise, the child is realizing that

[18] We'll later find (§5.7) that particulars also matter very much to how we should understand the relationship between perception and hallucination; in line with the emphasis on subjective indiscriminability in that context, however, their significance there is primarily phenomenological.

[19] A contrast familiarized by (Russell 1905; 1910-11).

before her discovery she was being *fooled*, but no fooling can have occurred if her thoughts were only ever descriptive.

Peter Strawson illustrates the shortcomings of purely descriptive modes of identification by appeal to the threat of "massive reduplication" (1959, 20). Suppose for *reductio* that our means of identifying particulars are entirely qualitative. No matter how thoroughly I describe the intrinsic properties of some persimmon I am seeing, it remains possible that there could be another persimmon somewhere in the universe that shares all the same properties. Nothing in the description settles which persimmon is the relevant one. What this shows is that descriptive conditions are insufficient to single out a concrete particular. But we *are* capable of identifying such particulars. So, contrary to the supposition, our means of identification cannot be purely descriptive.

It might be objected that some descriptive conditions do fit the bill: "the hamster she received on her eleventh birthday" is an example of the sort of description that could be satisfied by the original hamster but not by the replacement, let alone additional lookalikes in a far-flung region of the universe. This is a response that Strawson anticipates when he suggests that "[a]ll identifying descriptions of particulars may include, ultimately, a demonstrative element" (22). The proposed description is adequate because it exploits a demonstrative (non-descriptive) mode of identification; it still singles out the original hamster relationally, only it does so in a circuitous fashion. ("The hamster of *that child*" has taken over from "*that hamster*").

To streamline subsequent discussion, call the cases that Prior describes in terms of "thinking of Y where Y does not exist" *empty* cases. (This is somewhat of a misnomer by the lights of those for whom thinking of Pegasus is thinking about a special kind of object, but the label is also convenient and evocative.) Now notice that our need to go beyond the qualitative in our psychological ascriptions is not confined to cases where there really is an individual that the subject thinks about. It happens in the empty case; it also happens because of other confusions that induce failures of reference and aboutness, e.g., taking two things to be one.[20] If we care about individuals and not just types, it's to be expected that this concern will manifest itself when we are under the misapprehension that we have been dealing with an individual. Or consider the way that quantitative identity can come into our imaginative projects. I am imagining Pegasus; now I am imagining Pegasus transforming into a potato. This is different from thinking about Pegasus and thinking about a potato in immediate succession. But because of the radical qualitative changes involved in becoming a root vegetable, whatever purely descriptive condition I might plausibly have been using to single out Pegasus cannot explain my continuing to fix, in imagination, on a single entity. There is also a contrast between this

---

[20] Camp (2002) describes one such case, in which a person named "Fred" miscounts the inhabitants of his ant colony: "The ants arrive in two separate containers, one labeled 'little' and one labeled 'big.' Never reading the labels, Fred dumps the contents of both containers into the ant colony. He notices that there is a large number of little ants [...] and—he believes—one big ant, much bigger than the others. But Fred was inattentive and missed a second big ant when it tumbled into the ant colony. […] Fred's two big ants immediately arrange to split their time between running around on the surface performing feats of strength, and napping down in the bowels of the ant colony. They alternate up-time, one down below whenever the other is on the surface. As a result, even after Fred becomes fascinated by 'the' big ant, studying 'its' movements as often as he can, he never catches on that there are in fact two big ants." (27)

exercise and that of imagining *that there is a winged horse which is transforming into a potato*: the latter can be satisfied by imagining any number of individuals, but the former cannot.

These cases bring out an aspiration on the part of subjects, not just to an accounting of which properties are co-instantiated, but to thought about an individual—something singular and irreplaceable.[21] This is a feature of a subset of our intentional states, and it's what binds together *thinking of Y when Y exists* and *thinking of Y when Y does not exist* for the purposes of explanation.

The same observations should also reassure us that, contrary to the suggestion in §1.3.1, our expectation that some intentional states are relations between the thinker and what she thinks about is not just prompted by the syntax of the sentences we use to report these states. The source of the expectation, or anyway one clear source, is the conviction that we can think about individuals. Few would deny that some of our thoughts about things in the mind-independent world do indeed involve demonstrative modes of identification.[22] Prior's challenge is about the way that relations to particulars come into the equation: are these partly constitutive of the subject's non-descriptive thought, as *(a)* suggests ("thinking of Y *is* a relation"), or are they an additional element, extrinsic to the thought, that distinguishes the non-empty from the empty case? Part of what we shall discover, as we work towards a resolution of the challenge, is that the implications of these two answers differ more dramatically than they might seem to at first blush.

1.5 Three avenues

I'll now set out in more detail the three lines of response to Prior's challenge, which I'll label the *representational strategy*, the *ontological strategy*, and the *disjunctive strategy* respectively. Our inconsistent triad again, for ease of reference:

> *(a)* X's thinking of Y constitutes a relation between X and Y when Y exists, but *(b)* not when Y doesn't; but *(c)* X's thinking of Y is the same sort of thing whether Y exists or not. (Prior 1971, 130)

1.5.1 The representational strategy

The representational strategy is the product of rejecting *(a)*, or rather, a precisification of *(a)*. To achieve consistency, one needn't deny that, e.g., my thinking about the city of Toronto is a relation *at all*. When things go well, of course there winds up being a connection between the thinker's attitude and the object it is directed towards. The contentious question is whether my thinking of Toronto is *fundamentally* a relation between me and Toronto, hence a mental state that could not be if there were no such place. This is what the advocate of the representational strategy denies. If a relation plays a fundamental explanatory role here, according to the representational strategy, it's one

---

[21] In neutrally valenced senses of these terms. A piece of garbage is also singular and irreplaceable in the relevant sense—though the fact of something's being singular and irreplaceable takes on a new significance when that thing is valuable to us.

[22] Though this *has* been denied, notably by Russell. See discussions of the limited range of Russellian acquaintance, hence of demonstrative thought on his picture, in §3.1 and §4.1.1.

that relates a thinker to a mental representation: some element of her psychology that can, but need not, stand in a further relation to a *thing* (in the world) represented.[23]

There is a precursor of this approach in the proposal that the intentional object of an attitude is an idea, which appeals in a particular way to the subject's possession of an *idea of Pegasus* to explain her capacity to have Pegasus-thoughts. Anscombe (1965) identifies two immediate drawbacks of this position. One is that we already recognize *thinking of an idea of Pegasus* and *thinking of Pegasus* as two distinct activities, easily distinguishable just on an introspective basis (161). The "object as idea" proposal elides this contrast. A more severe version of this issue arises when we turn from the empty case to cases of thinking of what uncontroversially does exist beyond the mind, where the idea of the thing is now competing with the thing itself to be the object of the attitude. As Anscombe puts it, "one is thinking, surely of Winston Churchill, not of the idea of him, and just that fact started us off." (*ibid.*).

The lesson of this failure is that representations need to be invoked at the right point in the story for the representational strategy to be plausible at all. If a relation to a representation is essential to my thought about Toronto, it's not in virtue of its figuring as the object of the "thinking of" relation. Instead, the representation belongs in the account of what is going on with me subpersonally; my being in a certain representational state is what makes it the case that I am thinking about the city when this does occur. The main questions that are left open by this sketch of the structure of the representational strategy concern the nature of the representations that play the relevant role. These might be mental files,[24] name-like symbols in the language of thought,[25] the mental analogues of discourse referents,[26] or individual concepts,[27] for example.

1.5.2 The ontological strategy

The ontological strategy is the product of rejecting *(b)*, that is, the claim that X's thinking of Y is not a relation to Y when Y does not exist. Lest the position be manifestly inconsistent, the advocate of the ontological strategy needs to interpret the "exist" in this claim in a restricted fashion: there is no getting around the fact that a relation requires the existence of its relata. The idea behind the ontological strategy is that while the object of thought may not exist in one sense of that term, or exist in one *way*, it does have the ontological status necessary to be the second term of the proposed subject-object relation.

Historically this strategy has been associated with Brentano's student, Alexius Meinong. Meinong's version exploits a distinction between two modes of being, *existence* and *subsistence* (1960, 79). Pegasus, round squares, and golden mountains do not exist, but they do subsist, and this is supposed to explain how they are available to be thought about. Prominent contemporary implementations of the ontological strategy appeal to a different kind of object rather than a different kind of being. Possible worlds are called on to do a lot of philosophical work already; perhaps the objects of some of our

---

[23] Compare (Sainsbury 2018, chap. 6), where he proposes a response to Prior's challenge in this style. I return to that proposal in my own chapter 6 (§6.4).
[24] Recanati (2012). I discuss the mental-file framework in detail at §4.1.2.2.
[25] (Fodor 1975)
[26] (Kamp 1981)
[27] (Sainsbury 2005)

thoughts about non-existent (on this line of thought, non-actual) entities are possibilia.[28] It's also been suggested that the objects of these attitudes are fictions, abstract entities.[29] Advocates for these views tend to emphasize how minimal a demand the objects they invoke impose on the world—these are supposed to have the virtue of being ontologically "cheap."

1.5.3 The disjunctive strategy

Finally, we have the option of rejecting *(c)*, which says that X's thinking of Y is the same sort of thing whether Y exists or not. This approach is represented in the literature by the account of singular thought defended by John McDowell (1977, 1982b, 1987) and Gareth Evans (1982), about which I shall have much more to say below. For now I'll set out the disjunctive strategy in general terms.

What is to be gained by rejecting *(c)*? First, we can affirm that some intentional mental states are fundamentally relations to things in the world beyond our brains or bodies. Second, we can be faithful to folk ontology. Among the "objects" that we talk about, we habitually distinguish those that do from those that don't exist. We say things like, "Pegasus is not real," "There is no such thing as Pegasus," and so on. Unlike the ontological strategy (but like the representational strategy), the disjunctive strategy allows us to take these judgments at face value.[30]

A preliminary reason to be encouraged by this profile is that the tension that makes Prior's challenge challenging lies precisely in our being pulled one way and then the other on the relationality question, according to the kind of case we are considering. We feel that "thinking of" *must* be a relation in the most ordinary of cases; but also that it *can't possibly* be a relation in the empty case. The disjunctive strategy is the result of taking these two reactions very seriously, and, as a result, turning one's critical attention on the judgment of sameness of kind (represented by *(c)*) that pits them against each other.

Note, finally, that whereas the first two strategies provide the outline of a positive account of the empty case, the disjunctive strategy does not. It only settles what the empty case *isn't*, namely, a relation to the object of thought. There are, in fact, further constraints on an adequate spelling-out of the disjunctive strategy, which I discuss in chapter 3. For now, however, we can simply register how much is being left open by the disjunctive strategy relative to its competitors.

---

[28] See Lewis's (1978) account of names in fiction, as well as Priest's (2005) "noneism" and Berto's (2008; 2011) "modal Meinongianism."

[29] (Salmon 1998; 2002). This proposal should be distinguished from the much more widely-held view that fictional entities exist, the results of deliberate creative acts by novelists and so on—a view first proposed by Ingarden (1931). Salmon contrasts these claims in his discussion of two ways of handling the name "Vulcan" (2002, 101-103). (Note that these page numbers refer to the 2005 reprint).

[30] For a version of this point in the context of the representational strategy, see Sainsbury (2005, 66).

Chapter 2: The Task Ahead

I claim that the right response to Prior's challenge is to reject *(c)* or, in the terms of the last section, to adopt the disjunctive strategy. I don't just mean that the correct account of what it is to think about a particular thing is one that falls under this rubric, though of course I'm of the view that it does. I also mean that the only route to a genuinely *satisfying* response to the challenge—one that dissolves the tension we experience in trying to choose between the three claims—is by way of a disjunctive account of the mental states that *(c)* groups together. I expect this suggestion to be met with surprise, because while this approach is reasonably familiar, it is also poorly understood. My aim in the first part of this chapter is to remove some of the immediate obstacles to recognizing the appeal of the disjunctive strategy. In the second part, I set it in its wider context in the philosophy of mind.

2.1 Reframing the Challenge

While the disjunctive strategy affirms the first two of the three claims that Prior assembles, the standing of those claims differs. The claim that *X's thinking of Y where Y does not exist is not a relation to Y* is a fact to be accommodated. But the claim that *X's thinking of Y where Y does exist* is the disjunctive strategy's starting point, an expression of what motivates the approach. For this reason, the conflict between the disjunctive and the representational strategies—the latter being distinguished by its rejection of *(a)*—is particularly acute, and the representational strategy will be my main critical target in what follows.

A defender of that strategy might feel that we are forced by the slightest philosophical reflection into denying that thinking of an object is fundamentally a relation between it and the thinker. Acting on something, perceiving it—these require the participation of things in the world around us, just as a matter of definition. But the same cannot be true of our intentional mental states: it's just a datum that one can think about what doesn't exist. Surely this settles the question against the disjunctive strategy?

The problem with this line of argument is that what it presents as a "datum" is a conclusion drawn from a pair of observations that, while genuinely uncontroversial, fail to entail it. It's these observations that we are truly bound to in formulating a response to Prior's challenge, and recognizing as much already substantially improves the prospects for the disjunctive strategy. They are

> (i) that we speak of "thinking of," "wanting," "fearing," etc. whether or not we believe that the words that follow the attitude verb actually pick something out, and
>
> (ii) that it's not generally by means of introspection that one determines whether there in fact is such a thing as one's mental state purports to concern.

Let's consider each of these claims in more detail and determine what they do and do not establish.

2.1.1 Mental state ascriptions

I've already referred, in chapter 1, to G.E.M. Anscombe's (1965) comments on intentionality and perception.[1] I take Anscombe's claims to be an essential stepping-stone towards a correct appraisal of these issues, but they also have the potential to send us in the wrong direction. Here I'll explain what I understand her to be up to, and what I want to take away from her discussion.

Anscombe intervenes on what she sees as two versions of the same confused debate: one concerning the objects of thought, the other concerning the objects of perception.[2] In each case, the source of the confusion is a failure to recognize that certain transitive verbs permit of a use that others do not. Anscombe's corrective begins with a point about the meaning of the word "object," which she claims has undergone a historical transformation. On the "modern" sense, it's read as synonymous with "thing" or "entity" (158). So understood, talk of objects has existential import, and to say that a thought has an object is already to affirm *that there exists something* of which the subject thinks. But there is also an older philosophical usage on which "objects [are] always objects *of*—." (*ibid.*), so that talk of an object invites the question "object of what?" instead of triggering the expectation that there is a thing involved.

We are thus invited to consider whether certain philosophical problems might be generated by an assumption—too quick, unnoticed—about the meaning of words in our language. Taking "object" in the modern sense, the claim that every intentional state has an object entails that, for any state *I*, there exists some *x* such that *I* is about *x*. Consider a hiker who is, as we might say, *afraid of Bigfoot*. The object of her fear, the intentional object, is Bigfoot. This is a mythical creature: there is no entity we can point to as the target of the hiker's anxieties. The intentional object must then be some other kind of thing, perhaps an idea… And on we go down an unfortunate path.

The concept of an intentional object, Anscombe suggests, is not the concept of a kind of thing. It's a "grammatical notion" that helps us to describe and elucidate the workings of our language, just as the concept of a direct object does when we encounter it in elementary school (166). In fact, intentional objects are a "sub-class of direct objects" (162), characterized by a trio of special features. The object of an attitude may not exist; it may be indeterminate; furthermore, even if the expression in object place *does* refer to some entity, different ways of specifying that thing may not be equally suitable.

Recall Anscombe's case about the hunter who inadvertently shoots his father, and the distinction between "intentional" and "material" objects that she uses it to introduce (§1.3.3 above). The father in this context was the *material* object of the shooting. Material objects are once again not a type of thing, "what are now called 'material objects'—tables, planets, lumps of butter and so on" (168). This is another category of direct object. What the distinction between intentional and material objects tracks is a difference in the intuitive truth-conditions of sentences featuring attitude verbs. In determining whether a sentence with a material object is true, we are concerned with the existence of the entity indicated by the phrase that, as Anscombe would put it, "gives" the material object. (I.e., concerned with whether, in answering the question what the object

---

[1] (Anscombe 1965).

[2] The two debates *are* closely related and do call for similar treatments. I echo Anscombe's turn from thought to perception in chapter 5.

of the state or action was, we wind up referring to something.) When it comes to claims that feature intentional objects, we are concerned with giving voice to the subject's perspective. (I.e., concerned with whether, in answering the question about the object of the state or action, we say something that the subject would herself affirm in the moment.)[3]

The reader might be struck by the similarity between Anscombe's remarks and the claim that attitude verbs create intensional (and perhaps hyperintensional) contexts. Intensional contexts are marked by three features. One is the failure of existential generalization: that a subject *S* kicks an object *O* entails that *O* exists, for example, but that *S* paints a picture of *O* does not entail that *O* exists. A second feature is substitution failure: co-referring terms cannot be interchanged *salva veritate*. Suppose *S* fails to realize the identity of Hesperus and Phosphorus, and *S* believes that she is staring right at Hesperus. It may be false that *S* believes that she is staring right at Phosphorus.[4] A third feature of intensional contexts is that when certain expressions appear in the complement of the verb, these are ambiguous between specific and non-specific readings.[5] "Molly wants a burrito" or "Molly wants some burrito" may be read as attributing to Molly the desire for a particular burrito (specific reading), or attributing a desire that any burrito could satisfy (non-specific).

Is Anscombe just presenting us with a claim about the intensionality of attitude reports in different terms? Yes and no. On the one hand, the two descriptions do issue from and account for many of the same observations.[6] But the special value of Anscombe's description lies in the fact that it supplies a record of our linguistic behaviour that is both metaphysically and semantically neutral. With such a description

---

[3] In most ordinary cases, a single phrase provides both an intentional object and a material object. I am now thinking about the tree outside my window; there really is a tree of which I think (a "material" matter), and I think of it *as* a tree (an "intentional" matter). Even when the subject's perspective parts ways with reality because of confusion or ignorance, we can sometimes find a description that plays both roles. In the hunter example, for instance, we have: "that dark patch against the foliage" (167). While there is no stag, there really was a dark patch, so this intentional object is a material object as well.

[4] More precisely: in intensional contexts, substitution fails for expressions that are contingently co-extensive, while in hyperintensional contexts, this feature extends to necessarily co-referring expressions (like proper names, per (Kripke 1980). It's hard to deny that there are *intuitions* of substitution failure across the full range of referring expressions, but the verdict that propositional attitude reports are hyperintensional can nonetheless be resisted. Salmon (1986) and Soames (1987), for instance, trace these intuitions to a difference in the *aptness* of, e.g., "He believes that Hesperus is visible now" and "He believes that Phosphorus is visible now." Or one might follow (Braun 1998) in appealing to different ways of grasping a proposition.

[5] In Quine's (1956) original terminology, "relational" and "notional."

[6] The first and third of Anscombe's marks of intentionality map straightforwardly onto the first of the two features above, failure of existential generalization and failure of substitution. Her second mark is in the territory of the specific/non-specific ambiguity. But Anscombe blurs together, under the heading of "the object being indeterminate," the issue of whether a particular thing is in question and the extent to which an experience or representation is itself determinate (e.g. one can imagine a bird without imagining a bird with a particular number of feathers). Grouping these together is unproblematic in the present context, but the distinction between them will become important when we turn our attention to the perceptual verbs in chapter 5.

in hand, we have a new and a more expansive sense of what the possibilities are—of the various combinations of views about the mind and about language that could predict and explain these minimal facts about use. Those facts, in brief:

> *(i)* There is a core or "primary" use of intentional mental state ascriptions wherein the expression that follows the verb is intended both to pick out an entity that the subject thinks of (fears, wants, etc.) and to do so by means of a "description under which" (166) the subject herself thinks of that thing, where the measure of success for the utterance is the fulfillment of these two intentions.
>
> *(ii)* There is a secondary use of intentional mental state ascriptions (the "material" use) for which only the first of these intentions is in effect, and for which the measure of success is the fulfillment of this first intention alone.
>
> *(iii)* There is a secondary use of intentional mental state ascriptions (the "intentional" use) for which only the second of these intentions is in effect, and for which the measure of success is the fulfillment of this second intention alone.

This is what we need to hold on to as we formulate and evaluate hypotheses about the nature of the psychological facts that our ordinary mental state ascriptions are employed to express, and about the meanings of the terms that we express them with.

2.1.2 Subjective sameness

What's behind the impulse to affirm that thinking of Y when Y exists and thinking of Y when it doesn't is *the same sort of thing*? This claim poses a problem for the maintenance of *(a)* and *(b)* only if it is given a strict interpretation, as meaning that what we call "thinking of Y" in each case is a state of the same fundamental kind. Here there are two questions to consider. One is whether the *intuition* behind *(c)* either amounts to or implies a judgment about sameness of fundamental kind, or whether it instead issues from one of the various weaker commitments in the vicinity. The other is whether, on (philosophical) reflection, we find ourselves forced to adopt the claim in its strict version—or, if we are not *forced*, how severe are the costs we incur by rejecting it.

We already remove some of the trilemma-generating power of *(c)* when we approach the linguistic data with the attitude of caution I encouraged in the last sub-section. If one assumes that ordinary people speak univocally when they speak of *thinking of Y* across cases where Y does and where Y does not exist, then in insisting that these cases are different in kind after all one is going against the kind of common-sense understanding expressed in linguistic practice. What we have done with the aid of Anscombe's remarks on intentionality is retreat to a more conservative assessment of the evidence, leaving it an open question which philosophical picture will prove most consonant with our ordinary sense of things.

It remains a desideratum on an account of what it is to think about things in the world that it make sense of and even predict our speaking in the ways that we habitually do: it's plausible that if a philosophical position characterized such behaviour as

fundamentally confused, this would be a drawback to be compensated elsewhere (at least by the provision of an error theory). But now what we want to account for is just the fact that people speak or write a certain structurally-identical string of words across a pair of circumstances, and this observation can be accounted for in a wider range of ways than the more presumptuous "observation" that in so behaving people intend to be classifying two cases as cases of the same nature.

What might we be giving voice to, when we say that the non-empty and the empty case are the same sort of thing, if not the metaphysically significant judgment that in each the subject's mental state is of a single fundamental kind? A plausible basis for this claim is that we find the two cases to be *subjectively* the same: that whatever differences there are between the two circumstances are not revealed to introspection. This point is compatible with the truth of *(a)* and with the disjunctive strategy. Because of a common assumption about the significance of first-personal judgments in theorizing about the nature of the mind, however, one can slide imperceptibly from the metaphysically lightweight version of the claim to the metaphysically demanding one.

2.1.2.1 Introspection and infallibility

Ducks and decoy ducks, natural flavours and their artificial replicas, sincere and feigned enthusiasm, hunks of cheese and bars of soap—these are pairs whose members we may struggle, and sometimes fail, to distinguish.[7] They may be the same *from a given subject's point of view*, that is, the same according to her; but objectively they are not the same. They put up the same appearance, but they differ in reality. This is a distinction that, in matters concerning the outside world, we are accustomed to drawing.

In contrast, it is often suggested that there is no gap between appearance and reality within the mind; that the appearance *is* the reality. This is a claim that the disjunctive strategy requires us to reject—not across the board, but in some cases. We need to allow that questions about the subject's perspective and about her introspective judgments will receive different answers than questions about her mental state.

The equation of *being mental* with *being fully accessible to introspection* originates with Descartes, and emerges in a particularly striking fashion in his remarks on sensory experience. In the first Meditation, Descartes considers that what he previously took for a veridical perception of the room around him might be nothing more than a dream or a hallucination. He finds that he cannot rule out that he is in one of these delusive states, and so, as the method of doubt dictates, he casts aside those of his beliefs that rested on a perceptual foundation.

That is the first, negative step: the meditator's recognition of his ignorance about his objective situation, which leaves him feeling that he has "fallen suddenly into very deep water," unable to find either the bottom or the surface (AT IX 18, my translation). Descartes ultimately identifies a new foundation for his knowledge of the external world by arguing first that God exists and then that God is not a deceiver. But he also moves towards a more optimistic assessment of his epistemic situation in ways, which, because

[7] By way of evidence, I present a pandemic-era incident in which a person washed her hands for several days with what turned out to be a piece of cheddar cheese (u/lionellrichie 2020).

they don't rely on his failed argument for the existence of a benevolent God, are more consequential from the point of view of intellectual history:

> …it is I who have sensations, that is to say, who is aware of objects as though through the senses, since indeed I see light, I hear noise, I feel heat. — But all these objects are unreal, since I am dreaming. — Let it be so; certainly it seems to me that I see, I hear, and I feel heat. That cannot be false; that is what in me is properly called sensation; and in this precise sense, sensation is nothing but thought. (Descartes, AT VII:28/HR 1:153, trans. Williams 1978, 63)

Perception cannot be counted on to play its usual epistemic role, that of furnishing us with knowledge of the external world. But in any case of (putative) perception, Descartes suggests, something *is* known: the experience itself, the sensory appearances, the way things seem to him. As Bernard Williams (2005) explains, Descartes begins in this passage to classify as thoughts (*cogitationes*) "the purely mental element in experiences which earlier he treated as presupposing the existence of his body and the physical world" (64). Having previously cast doubt on experience as a whole, "[n]ow he is prepared to 'shear off' a purely mental experience, and call that 'sensation'" (*ibid.*).

Myles Burnyeat identifies this "shearing off" as the move that sets modern scepticism apart from its precursors—isolating the subject from the whole of the external world, her own body now included in that category. It's at this point that Descartes transitions from "[focusing] in a largely traditional manner on what is true or false of the real world" to "[speaking] of things being true or false without meaning true or false of a real objective world outside the mind." (1982, 38) He "discovers . . . the truth of statements describing the subjective states involved in the process of doubt itself." (*ibid*.)

Why does this move distinctively entrain, as John McDowell (1987, 238) likes to put it, "the loss of the world"? Descartes's new picture looks initially to set a limit on scepticism: even in the midst of the most radical doubt, there is still something that we know (how things seem to me, that I am thinking). Picking up where Burnyeat leaves off, however, McDowell emphasizes that there is a relationship between access to the external world and certainty about one's own mind:

> In a fully Cartesian picture, the inner life takes place in an autonomous realm, transparent to the introspective awareness of its subject; the access of subjectivity to the rest of the world becomes correspondingly problematic. (236)

In order for it to be plausible that we are infallible judges of our own mental states, the mental needs to be construed as wholly and unproblematically discoverable by us. But clearly we don't have that kind of access to any part of the mind-independent world. Our judgments about the mind can only have the security Descartes ascribes to them, then, if the mental is disconnected from the rest of the world; it needs to be self-contained (hence McDowell's talk of an "autonomous realm"). But if the mind and the world are disconnected in this way, then we're in a position of radical uncertainty about the latter.

McDowell makes this observation to bring it to attention that Descartes' trade-off is not the only way of going here. If we accept that our insight into our own minds is of the ordinary, reliable-but-fallible variety, that's consistent with holding that the subjective and objective are integrated: that "objects themselves can *figure in* thoughts that are

among the contents of the mind" (237; emphasis original). This is the trade-off characteristic of the disjunctive strategy.

Of course, Descartes doesn't pull the idea that we are infallible about the nature our own mental states out of thin air; he doesn't *decide* to trade access to the external world for certainty about a realm of subjective facts. Instead, he takes his sceptical exercise to reveal that our epistemic situation is like that. We've yet to address the evidence for that conclusion. On Descartes's telling, there is a subjective fact that I have knowledge of even if I am, unbeknownst to me, being deceived by an evil demon. As the foregoing has helped to bring out, he took a step to get to that claim (the "shearing off"). A step from where? What's clear is that there are claims we are licensed in making however uncertain or confused we might be about the situation we are in; claims that, so long as we utter them sincerely, are bound to be true (or, even more conservatively, treated as true) in virtue of that very fact. This is the feature that philosophers often express by saying that some first-person statements are *incorrigible*.

The claim that we have special epistemic access to our own mental states is one way of spelling out this incorrigibility, but there are other, more deflationary treatments of it too. We've noted that the Cartesian gives a *positive* account of incorrigibility: there is a fact that the subject knows, and about which she is not fallible.[8] On this story, the way that we handle first-person reports is a consequence of our recognition of an independent epistemic truth. Against Descartes it has been suggested that "the 'incorrigibility' of descriptions of sense-contents (narrowly understood) is not any kind of correctness."[9] It's possible, for instance, that our treating these reports in the way that we do is somehow constitutive of the reality that the Cartesian then interprets through the framework of knowing a truth. Giving this claim the extensive discussion it deserves would take me beyond the scope of the present project, but we can work with the following thought: that when we decide to narrow in on a subject's point of view, we are expressly *not* attending to what is objectively the case. This *disregard for what is objective* or third-personally available is portrayed by the Cartesian as *regard for each person's special knowledge of their own inner world*. The possibility of accounting for incorrigibility in another way shows that it is not obvious that our epistemic situation vis-à-vis our own minds is as Descartes claims.

2.2 The requisite structure

The task of construing a kind of mental state as fundamentally relational is a delicate one. We need to combat the temptation to decompose states of that kind into a "mental" and a "physical" element right off the bat; but we also run the risk of characterizing the target state in such a way that, contrary to our intentions, it proves to

---

[8] Burnyeat (1982) declares that with Descartes "[s]ubjective truth has arrived to stay, constituting one's own experience as an object for description like any other." (38-40) But these are not "mere" subjective truths. As Williams explains, for "I think, I exist" to do the work that Descartes wants it to, it must express the meditator's recognition, from within the subjective perspective, of an objective fact (1978, 84). "According to myself I do exist" is not an interesting philosophical claim. "I think, therefore I *am*" is, because it is a claim about what there is in reality.

[9] (Anscombe 1981, 54). Compare (Wittgenstein 1953, §§288–289). For a recent case against accounts of incorrigibility in terms of a privileged epistemic relation, see (Bar-On 2004).

be decomposable after all. The task and the difficulties just cited are not unique to the case of singular thought. I'll first discuss their expression across a range of contexts, and then begin filling out the structure of the view that I want to defend under the heading of the disjunctive strategy. The central concern will be how to handle the relationship between singular thought about an existing thing and whatever is going on in the empty case.

2.2.1 Resisting analysis

Say we set out to vindicate the idea that for a subject to be in a kind of mental state *M* requires not only that her brain and body be a certain way, but also that the wider world be a certain way. Suppose that we also

> (i) acknowledge the possibility of a kind of mental state N such that when S is in an N state it is from her point of view as if she is in an M state, though the environmental circumstance partly constitutive of M does not obtain, and
>
> (ii) want to deny that M is adequately explained in terms of the conjunction of N and aforementioned environmental circumstance.

What we are after in this case, in Timothy Williamson's (1998) terminology, is an account of *M* as both "broad" and "prime" (390, 392). Williamson introduces the notion of a *broad* condition as follows:

> A case α is *internally like* a case β iff the total internal physical state of the agent in α is exactly the same as the total internal physical state of the agent in β. A condition C is *narrow* iff for all cases α and β, if α is internally like β then C obtains in α iff C obtains in β. In other terminology, narrow conditions supervene on or are determined by internal physical state; no difference in whether they obtain without a difference in that state. C is *broad* iff it is not narrow. (390)[10]

The existence of broad conditions is neither controversial nor at odds with internalism about the mental, the view that a subject's mental states supervene on her intrinsic physical properties. What it takes for a concept to be a concept of a broad condition is just that there are requirements on a correct application of that concept that take us beyond the total internal physical state of the subject (the internalist will accordingly deny that concepts like *knowledge* and *perception* are concepts of mental states).[11] As Williamson emphasizes, many of our ordinary mental state ascriptions are broad (390–391). The interesting question is whether some of the broad conditions are also *prime*:

> A condition C is *composite* iff it is the conjunction of some narrow condition D with some environmental condition E. As a special case, a narrow mental condition is trivially composite, for it is the conjunction of itself with the

[10] Williamson is assuming for argumentative purposes in the cited paper that "the total internal physical state of the subject and the total physical state of the external environment jointly determine the total state of the world" (1998, 390)—that the psychological facts, whether narrow or broad, are determined by the physical facts.

[11] Compare (Williamson 1995, 533).

> environmental condition that holds in all cases whatsoever. C is *prime* iff it is not composite. (392)

To say that thoughts about particulars are relations to those particulars, in the sense of that claim that the disjunctive strategy affirms and the representational strategy rejects, is to say that these mental states are prime.

As the foregoing suggests, a first step towards an understanding of *M* states as fundamentally world-involving is to insist that the concept of *M* is the primitive notion. Williamson's "knowledge first" slogan is perhaps the best-known expression of this move. But an earlier and a more general instance, which will help us get a feel for its dialectical impact, appears in the third chapter of P.F. Strawson's *Individuals* (1959). In that chapter, Strawson takes issue with a pair of views about the nature of persons. These views purport to be in opposition, and in some respects they are; according to Strawson, however, they also have a basic error in common. On the *Cartesian* view, a person is a combination of two more fundamental entities, a body and a mind. On the *no-ownership* view, which is inspired by some claims of Wittgenstein's, the Cartesian errs by positing an ego to which experiences are ascribed. Experiences are "causally dependent on the state of a particular body," but this is not a relation that should be cashed out in terms of ownership, and we certainly should not introduce an immaterial mind to serve as the owner (96).[12]

As Strawson characterizes the no-ownership view, it is insufficiently radical as a response to Cartesianism. It goes along with the division of a person into two components and then rebels by dispensing with the more dubious of these, the ego. Whether or not this is a faithful characterization of anybody's view (Strawson is explicit that his aim is to make a point rather than target a particular opponent), the interesting charge that Strawson uses it to articulate is that once the Cartesian division is made there is no hope of producing a plausible account of the concept that we set out to elucidate. The only escape from the problems that Strawson takes to beset both the Cartesian and the no-ownership approaches, he suggests, is to "acknowledge […] the primitiveness of the concept of a person" (101). About that concept, he has the following to say:

> What I mean by the concept of a person is the concept of a type of entity such that *both* predicates ascribing states of consciousness *and* predicates ascribing corporeal characteristics, a physical situation &c. are equally applicable to a single individual of that single type. (101–2)

Our sense that there is a clear line to be drawn between two types of predicates that we apply to human beings—predicates like "is in pain," "is grumpy," or "is calculating the tip," on the one hand, and "is 5 feet tall," "is from Eritria," or "is wearing a red sweater," on the other—is a driver of our efforts to analyze the concept "person" into a mental and a physical component. This, and not whether we go on to make the mental its own kind of substance, is the fundamental issue.

---

[12] See (Anscombe 1975) for a view that belongs in the "no-ownership" category, at least according to this minimal description. Whether Anscombe has the sorts of commitments that Strawson latches onto in his critique of that position is a further question.

Strawson is here advocating for a certain kind of solution to a certain kind of problem in the philosophy of mind. The problem is the failure of a foundational idea or a framing assumption (it's important, in this connection, that Strawson interprets the no-ownership view as itself a legacy of Cartesianism), namely that various of our ordinary concepts are appropriately analyzed into "internal" and "external," "mental" and "physical" components.[13] This failure is revealed when we encounter difficulties that are insoluble from within the received framework: in the case that Strawson is concerned with, the main symptom is an "[oscillation] between philosophical skepticism and philosophical behaviourism." (109) The solution is to treat the contested concept, the would-be analysand, as *primitive*. In the notion of primitiveness there is at once the suggestion that the target concept is in one or more senses prior to the proposed components, and also that, at least by the standards of some conceptions of what it is to give a philosophical explanation of a phenomenon, it is not amenable to explanation at all. In particular, it is not amenable to the kind of explanation that will decompose it into what are supposed to be more tractable units, or take the air of mystery dispersed across a field of inquiry and restrict it to a smaller region (where the latter charge might be leveled against substance dualists in particular).

Strawson's remarks on the concept of a person exemplify in its most general form an approach that one can take to any number of associated concepts, concepts of states of persons, capacities that they have, and activities that they engage in. Strawson himself also discusses predicates like *being depressed*, *smiling*, and *going for a walk*, which his opponents would carve up along the same mental-physical lines. In epistemology, this style of response yields a knowledge-first account of the relationship between knowing and believing (Williamson 2002). In the philosophy of action, we get the proposal that intentional action is basic, not something that can be ginned up by the addition of a special mental feature to a case of bodily movement.[14] In the philosophy of perception, it manifests in the view that perception is fundamentally a relation between the subject and her environment.[15] As applied to the problem of other minds, the strategy is to conceive of behaviour as an expression of—rather than as a bodily occurrence that is merely *caused* by—the subject's thoughts and feelings (McDowell 1982a).

What these views all have in common is that they put the success, the "good case," the broad condition *first*, treating it as explanatorily basic and conceptually primitive. The Cartesian alternative is to fix on something that is common between successes and failures, which is (by Cartesian lights) strictly mental or strictly physical, and make that the basic building block of the phenomenon that bridges the mind-world gap. [16] Once we've resolved to treat the concept of knowledge, of depression, of

---

[13] This decomposition is, in Descartes' case, epistemologically motivated. It gains new steam when the method of analysis becomes an object of philosophical enthusiasm in its own right.
[14] Anscombe appears to have held this view (1957: §19). See also Wilson (1989), Ford (2011), Levy (2013), and O'Brien (2017) among others.
[15] See Campbell (2002), Martin (2004), and Brewer (2011).
[16] Whether this building block is on the mental or physical side of the divide depends on the point of view that the subject engaged in sceptical reflection occupies on the target mind-world phenomenon. As concerns my perceptual access to the external world, the question that skepticism presses on me is whether I can get beyond my own subjectivity to something objective—I face the threat of being stuck in my first-person point of view. What I seem to have

perception, of behaviour, etc., as primitive, the next question is what to make of the supposed common element that excited the Cartesian. This is where the disjunctive strategy comes into play, but it is not the only option, nor are the various contexts where it might be applied perfectly analogous.

2.2.2 Peaceful coexistence?

We've now noticed that when one considers the sorts of states, capacities, and events that we conceive of as requiring more of the world than merely that a subject's brain and body be some way, one often finds waiting in the wings something that might be identified as the more basic, "strictly mental" component of the world-involving phenomenon. Alongside knowledge, for example, there is belief. Alongside perception, there is sensory experience, or how things sensibly appear. To insist on the primitiveness of knowledge or perception, one needs to replace the tempting, "common-factor" interpretation of these nearby concepts with an alternative. I'll first present a non-disjunctive account of one of these putative common factors and why I won't be taking this approach in the context of singular thought. I'll then turn to disjunctive proposals.

Timothy Williamson (1995) claims that knowledge is *both* factive *and* a mental state: his position accordingly constitutes a rejection of (wholesale) internalism about the mental. He emphasizes that the main point of knowledge-first epistemology is to refuse "conjunctive" accounts of knowledge, that is, accounts of it in terms of belief (an internal state of the subject) and various further conditions (at least one of which, truth, depends on the state of the wider world). He claims that this position is "neutral on whether believing is necessary for knowing" (534), i.e., compatible with an affirmative answer to the question about the necessity of belief. That's possible because *A* can be a necessary condition for *B* without being a *part* of *B*.

How does one resist a characterization of belief as the "highest common factor" between knowing and mere believing if not by denying that a single kind of state is present both when one knows and when one merely believes? What Williamson denies is that believing is "a factor of knowing at all (whether or not [it] is a necessary condition)" (1995, 563). Going from merely believing to knowing is not a matter of some additional conditions supplementing one's belief, such that one's mental state itself remains fundamentally unchanged; it's "to be in a new kind of mental state, whose essence involves the world" (563). This can be so whether because the new mental state replaces, or because it simply coexists with, the state of believing.

Williamson is in favour of pursuing this line of response elsewhere, e.g., as concerns the relationship between perception and how things sensibly seem. As a

---

in hand regardless, hence the candidate building-block of the success condition, is something mental: the sensory appearances, a belief, etc. These become perception and knowledge with the *addition* of something non-mental. As concerns my access to the mind of another person, in contrast, the worry is about getting from the objective (what is going on with a body) to the subjective (there being thoughts and feelings behind the movements). In this case, what I have in hand from the sceptical point of view is, say, *smiling conceived merely as a way that a face can be positioned*. The good case, in which a smile is occurring *because* the relevant person is happy, is again to be understood conjunctively, but in the opposite direction from the perceptual or epistemic examples: it's smiling "plus" (caused by) a feeling of happiness.

preliminary to considering such an extension, I want to note some differences in the folk-psychological lay of the land across the contexts where it might be made. In the epistemological context we have the concepts of *knowledge* and *belief*, the former being a broad and the latter a narrow condition. In the perceptual context we have perception, hallucination, and illusion (we do talk in everyday life about how things appear, but it's not clear that there is a pre-existing *concept* of 'the sensory appearances,' with a role analogous to that of belief, that philosophers have simply come across). But the concepts of hallucination and illusion, like the concept of perception, invoke features of a situation beyond the internal physical state of a subject. How they contrast is that in the case of delusive experiences the further requirements are negative in character, so there is not *one specific way* that the world needs to be in order for the experience to be an experience of that type: the world is *not* as it appears to the subject, the sensory state is *not* caused by the normal route for perceptions.

In the case of the mental states primarily at issue in this dissertation, things are different still. We have such notions as *thinking of*, *loving*, *wanting*—and that's it. There is not an obvious distinction here among concepts that bring in the state of the world and concepts that don't; instead, there is a strong temptation to think that all of these concepts are concepts of narrow conditions, though our ascriptions of intentional mental states themselves involve a broad condition when the object of the attitude is a constituent of mind-independent reality.[17]

My contention is that the "attitude" and "object" parts of our intentional states cannot be hewn apart in this way: *thinking of something* is the basic notion, and it is a relation between the subject and that "something." I began to suggest in the last chapter (§2.1.1) that the cases where we say that a subject thinks of something are of two different kinds, and that our utterances may themselves be of a correspondingly different nature. Were we to try and characterize this situation non-disjunctively (while still insisting on the primitiveness of the broad condition of having a singular thought) it might look something like this:

> *Bad (empty) case:* the subject has a thought that is singular by some intrinsic measure.
>
> *Good case:* the subject has a thought that is singular by some intrinsic measure, and she also has a thought that is "extrinsically" singular (irreducibly relational).[18]

---

[17] To elaborate: the idea is that there is one kind of mental state—thinking of, wanting, fearing, etc.—where the "having an object" characteristic of these states is an intrinsic feature of them (e.g., a matter of their possession of a representational content). If we add that the object of thought is something out in the world, then we plausibly have a broad condition, which, however, advocates of the representational strategy would deny is prime. The closest analogy here is with certain uses of the perceptual verbs, as in, for instance, "Macbeth sees a dagger," which I discuss in chapter 5.

[18] Williamson's setup, for comparison: in the bad case the subject believes, in the good case the subject believes and knows.

This account of the good case might appear redundant, but the same could be said about combining belief that *p* and knowledge that *p*, so let's pursue it a bit further.

M.G.F. Martin (2006) has suggested that what makes Williamson's approach suitable for the case of knowledge and belief is that these states plausibly have distinct explanatory roles. The same may be true of perception (the success notion) and sensory experience (the putative common factor) relative to some questions, but Martin claims that "there is also a central area of concern where the two notions are bound to be in competition" (371). That area of concern is the subject's sensory phenomenology, how things are for her visually, auditorily, etc. at a given time. If *both* a perception *and* a sensory experience were the sorts of things that could figure in the subject's stream of consciousness, then they might be allowed to co-exist in the way that Williamson allows for belief and knowledge. Yet it seems that they cannot:

> …if we are intending to give the most determinate and specific account of how things are with a subject at a given time, thereby picking out their conscious state, either what we pick out is the perceptual episode, as the naïve realist supposes, or the common sensory experience, as defenders of the Common Kind Assumption suppose. (*ibid.*)[19]

If this framing is correct, then determining whether we have an alternative to a disjunctive account of *thinking of something* is a matter of determining whether there is a similar sort of "explanatory exclusion" issue as concerns singular thought.

I suspect that there is. A first, albeit defeasible, indication in this direction: as above, in neither the perceptual case nor the objectual attitude case is there an ordinary concept with a status analogous to the one that belief has in the epistemological context—a concept of what is *both* uncontroversially a narrow condition *and* plausibly a common feature of the "good" and "bad" cases. One obvious explanation of the fact that we have a concept of belief *and* a concept of knowledge is that we have two distinct explanatory roles in view. Likewise, perception and hallucination are two concepts of sensory states that are *not* in the kind of competition that Martin describes as occurring between the concept of perception and the philosopher's particular concept of sensory experience (as a common element between perception and hallucination, or a broader category that both fall into). With objectual attitudes like "thinking of," "wanting," etc., we just have the one concept, whose status as the concept of a prime or of fundamentally a narrow condition is itself contested. This suggests that there is a *one* prominent or central explanatory role here, and that "thinking of" qua internal state and qua relation to a thing are both in competition to play it.

One of the core issues in the competition between the representational and the disjunctive responses to Prior's challenge is whether anything is lost—from the perspective of psychological explanation, and relative to our account of the subjective point of view—if what we cite is not a relation between the subject and the object of her thought, but an intrinsic state of the subject. The relation and the intrinsic state may not immediately exclude each other, in the way that perception and sensory experience exclude each other because there is only "space for" one of these in the subject's stream of consciousness at a time. But they nonetheless are called on to do the very same work.

---

[19] Similar considerations figure in (Martin 2004, 59), though not as addressed to Williamson.

In particular, we might understand the situation as follows. The narrow state is introduced to take over from the broad one for the purposes of explanation and description because we find ourselves wanting to make the same explanations and descriptions in the empty case, where the broad condition doesn't obtain—because, in other words, of the apparent impossibility of accounting for the empty case by reference to the broad condition.[20] But the account I will defend in this dissertation provides us, as I explain in the next chapter, with a way of doing just that. From this point of view, the introduction of the narrow state was ill-motivated. So while a Williamson-style, non-disjunctive account may not strictly be ruled out, the inclusion in our story of what is in effect an idle piece of machinery has nothing to recommend it.[21]

2.2.3 Disjunctivism

Disjunctivism originated as a position in the philosophy of perception. In perception, one succeeds in making contact with the world; in hallucination, one fails to make contact with the world. A hallucination may be subjectively indiscriminable from a given perception; the same description of the sensory appearances seems to be called for in both cases. These observations have led many to conclude that perception and hallucination are of a single experiential kind, locating the contrast between them beyond the mind of the subject. Such is the response characteristic of so-called "common factor" accounts of the metaphysics of perception, notably sense-datum theory and intentionalism.

Traditional sense-datum theorists like Moore and Russell held that in perception and hallucination alike the subject's experience was accounted for by her awareness of sense-data, conceived as both mind-independent and non-environmental (i.e., private).[22] Later sense-datum theorists give the same role to a relation to sense-data construed as mind-*dependent*.[23] According to intentionalists, sensory experience is representational in nature. Whether one perceives or hallucinates, one is in a state that represents one's environment as being a certain way; veridical perceptions are distinguished in being *accurate* representations as well as by their etiology.[24]

---

[20] We may then resolve our introduction of the narrow state with our antecedent commitment to the broad condition by positing that the former is a component of the latter.

[21] This is not to deny that some narrow state could both coexist with singular thought (qua broad) and play a quite different explanatory role; see the introduction of Evans's "informational states" at §3.3 below. What I intend to rule out is the possibility of such a narrow state playing a role roughly analogous to that played by belief in the absence of knowledge, i.e., being independently sufficient to explain what we want to explain about the empty case.

[22] Though note that when Moore first entertains the idea of sense-data, he is open to the possibility that they can be identical with the surfaces of ordinary things. See his discussion of an argument against this proposal in (1918–19, 18–25); he revisits similar issues in (1925, 219–222). Moore he moves towards the proposal that the relation between sense-data and what they are of is always representational, in his (1952). For an early statement of Russell's view, see (Russell 1912, chap. 1).

[23] (Jackson 1977), (Robinson 1994).

[24] Anscombe (1965) is often cited as a source of inspiration for this view; Armstrong (1968) and Pitcher (1971) argue explicitly that perception is a form of belief, which is one way of classifying sensory experience among the intentional states. Prominent statements of intentionalism come from

Against this, naïve realists hold that perception is fundamentally a relation between the subject and her environment; elements of the perceived scene are not just represented in, but are constituents of, her experience.[25] These conditions do not obtain in a case of hallucination; hence the pressure towards a disjunctive account of sensory experience and of the sensible appearances. The availability of such a position was first brought to attention by J.M. Hinton (1967b). Hinton observes that

> "It is to me as if I saw a flash of light." (217)

might be used to express *either* that one saw a flash of light *or* that one suffered an illusion of a flash of light. Paul Snowdown spells out the proposal in more general terms as follows:

> The disjunctive picture divides what makes 'looks'-ascriptions true into two classes. In cases where there is no sighting [i.e., where the subject is not veridically perceiving] they are made true by a state of affairs intrinsically independent of surrounding objects; but in cases of sightings the truth-conferring state of affairs involves the surrounding objects. (1981, 186)

Nothing is said, in either of these extracts, about what precisely is going on when one suffers an illusion or hallucination as of, say, a flash of light. The main point is to indicate that when we talk about experiences we might sometimes be talking about perception and sometimes not—rather than always talking about a single phenomenon, sensory experience, that occurs both when one perceives and when one hallucinates.

There is a distinction within disjunctivism between *positive* and *negative* varieties. These are two ways of instituting the idea that veridical perceptions, and those sensory experiences that fall short of them, are states of fundamentally different kinds. The labels "positive" and "negative" refer to the way that these two varieties of disjunctivism handle non-veridical sensory experiences: by supplying a positive account, or by characterizing them purely negatively.

Consider a version of the argument from illusion, as traditionally offered in support of sense-datum theory:

> (1) In an illusion, object *O* seems *F* to me. (E.g., a stick looks bent.)
> (2) If it sensibly appears to me that something is *F*, there must be something I am aware of which is *F*.[26]
> (3) *O* is not *F*. [Such is the nature of illusion.]
> (4) Therefore, there must be some object $O' \neq O$ which I am aware of and which is *F*. Call this further object (*O'*) a "sense-datum."
> (5) There should be a single account of the subject's experience in both illusion and veridical perception.

Peacocke (1983), Searle (1983, chap. 2), Harman (1990), Tye (1992), and Dretske (1981; 1995); recent defenders include Byrne (2001), Pautz (2007), and Siegel (2010).

[25] Campbell (2002), Martin (2002b; 2004), Brewer (2011). See also (Fish 2009) for critical engagement with the view.

[26] This is what Howard Robinson terms the "phenomenal principle" (1994, 32).

(6) Therefore, in cases of veridical perception and illusion alike, the immediate objects of awareness are sense-data.

The characteristic disjunctivist move is to resist the so-called "generalizing step" of premise (5), which is usually supported by claims about the subjective resemblance of perception and non-veridical experience.[27] Rejecting this step is the minimal intervention required to prevent the triumph of a common-factor view, and as applied to the above argument, it yields a positive disjunctivism: in veridical perception one is immediately aware of ordinary things, whereas in non-veridical sensory experiences one is immediately aware of sense-data.

This reactive proposal illustrates a challenge for disjunctivism that might lead us away from pursuing it in a positive form. One of the facts about sensory experience that any theory is accountable to is that we can be taken in by hallucinations and illusions, mistaking them for veridical perceptions. The proposal that the subject's phenomenology is explained in one case by her unmediated awareness of things in the environment and in the other by her unmediated awareness of sense-data seems to make it more puzzling, not less, why experiences of the two kinds should strike us as phenomenologically the same—and likewise for other ways of characterizing non-veridical experiences in positive terms. In this way among others, the approach might strike one as ad hoc.[28] This is a first hint that the explanatory potential of disjunctivism may lie elsewhere than in the provision of a distinct account of non-veridical experience.

` Then there is the pressure exerted on the disjunctivist by the following line of thought (characteristic of *causal* versions of the variously construed "argument from hallucination").[29] While the distal cause of a given veridical perception is the thing in the environment that the subject is aware of, the proximate cause is a pattern of neural activity. Take a token visual experience *E*. Suppose that your brain is artificially stimulated to reproduce the patterns of neural activity realized when you undergo *E*. By the principle "same proximate cause, same immediate effect," the resulting experience *E'*

---

[27] In the early sense-datum theorists' formulations of the argument from illusion, this often took the form of an observation about the continuity of normal and illusory conditions. See for instance (Price 1932, 32), (Ayer 1940, 8–9), and Broad (1952, 9). For exposition and discussion of several different arguments towards the sense-datum theorist's desired conclusion, see (Pitcher 1971, 13–50).

[28] One way to address this concern is to finesse what I have thus far presented as a bare assertion of metaphysical difference, though the view that results may no longer be strictly disjunctive in character. For example, Sethi (2020) argues that in perception we are acquainted with mind-independent instances of sensible qualities (borne by ordinary things in the world) while in hallucination we are acquainted with mind-dependent instances. Sethi makes two key moves. The first is to propose that sensible qualities are *ontologically flexible* (9), so that, contrary to the simple fusion of naïve realism and sense-datum theory I have sketched above, the account of the subject's phenomenology in perception and hallucination is the very same (awareness of instances of sensible qualities). The other is to argue that the sensible instances we are aware of in perception are *over-determined*: guaranteed both by the perceiver's internal state and by the object she perceives (12–17). This last move is what ensures that the mind-independent instance is not redundant in the explanation of the perceiving subject's phenomenology, which prevents the view from collapsing into a version of sense-datum theory.

[29] Recall our brief encounter with this version of the argument in §1.3.2 above.

will be qualitatively indistinguishable from *E*.[30] This suggests that being in a certain brain state is sufficient for the occurrence of a sensory experience.

The causal argument provides us with a new, non-phenomenological basis for the generalizing step I cited above, the thought being that if a perception and a hallucination have the same proximate cause they are experiences of the same kind. The disjunctivist can intervene in similar fashion her, for instance by emphasizing the role of objects as constituents and not merely as causes of veridical perceptions (this amounts to suggesting that experiential kinds are not individuated purely causally). Yet there is a more persistent challenge here that this first line of response does not address.

It's just a stipulation of the case we considered that the kind of brain state responsible for our imagined hallucination also occurs in the corresponding perception, and it's hard to deny that this could be so. Furthermore, there are numerous similarities across cases of perception and hallucination that need to be accounted for, like their subjective indiscriminability and the reactions and behaviours that both can bring about. These are explicable in a case of perception by reference to the subject's awareness of things in her environment. But the causal argument suggests that they can be explained across a *wider* range of conditions—perception and hallucination alike—by reference to the subject's brain state. In this way, awareness of the mind-independent world is "screened off" (Martin 2004, 62) from playing the explanatory role that the naïve realist wanted to accord to it, and we find ourselves led back to a conjunctive account of perception: the subject's experience is wholly accounted for by her brain state, whose causal history is what qualifies the experience as perceptual rather than hallucinatory.

M.G.F. Martin's formulation of negative disjunctivism is responsive to this last challenge.[31] Martin proposes that in the case of causally matching hallucinations like the one introduced above, "there is no more to the phenomenal character of such experiences than that of being indiscriminable from corresponding visual perceptions" (2006, 370). How does this help to address the screening off problem? The negative epistemic property of *being indiscriminable from a veridical perception* is one that both hallucinations and perceptions possess (the latter have it trivially). It thus has the same generality as the brain-state explanation that the causal argument introduced. Yet it achieves this generality without putting hallucination and perception on an equal footing: the character of the former is entirely parasitic on that of the latter. For this reason, veridical perception retains its distinct explanatory role.[32]

Abstracting from the details of Martin's view, we can take notice of a broader strategy that it instantiates: that of treating the two disjuncts *asymmetrically*. The negative disjunctivist's account of hallucination makes essential reference to veridical perception, and not vice-versa. In contrast, the common-factor theory provides a specification of the subject's experience that is neutral between perception and hallucination. The

---

[30] See, e.g., (Robinson 1994, chap. 6, sec. 2).

[31] This should not be taken to imply that the whole point of the position is to avoid a trap laid by opponents of naïve realism. That it avoids the challenge in question here is a prominent virtue, but the view is also motivated by considerations about the phenomenology of sensory experience. The latter are a central theme in (Martin 2002b, sec. 2).

[32] Naturally this brief survey skims over some of the subtleties in the area. For a more comprehensive explanation of the appeal to indiscriminability as a response to the screening-off problem in particular, see (Martin 2004, 69–70).

asymmetrical structure of Martin's negative disjunctivism is shared by the account that I will introduce in the next chapter. But the dependence of one thing on another can take various forms, and this will be a point of difference between the approaches.

On the view I develop, the emphasis is on the dependence of pretence on reality, rather than the dependence of what is merely indiscriminable from a token state on the token state itself. The view we shall arrive at by pursuing this line does not fit neatly into either of the familiar kinds of disjunctivism, positive and negative. On the one hand, *no positive element* of the subject's psychology will be posited in the empty case which is sufficient to make sense of her subjective life or her behaviour. On the other hand, there will be a positive *account* of the empty case, and in this way the view contrasts with negative disjunctivism. Part of the interest of the proposal thus lies in the fact that it represents a novel way of threading the needle required to avoid falling into a conjunctive account of the target mental state.

Chapter 3: Minds and Make-Believe

In the previous chapter, I set the disjunctive response to Prior's challenge alongside views of a similar spirit in epistemology, philosophy of perception, and philosophy of action. I then described, still in quite general terms, the structure that our implementation of the disjunctive strategy would take. In this chapter, I introduce the account of singular thought that fits this profile, due to John McDowell and Gareth Evans. My primary aim is to draw attention to an underexplored way of handling empty cases and the various difficulties that such cases apparently pose for an advocate of the disjunctive strategy. §§3.1–3.5 I spell out the genesis and content of this proposal, drawing primarily from (Evans 1982, chap. 10). In §3.6 I use it to respond to a series of pressing objections against accounts of singular thought as object-dependent, hence also against our pursuit of the disjunctive strategy; this will help to clarify what is distinctive about the view and provide a first indication of the fruits it bears.

3.1 No object, no thought

In *The Varieties of Reference*, Evans defends a neo-Fregean, two-level account of meaning: the contribution of a singular term to a sentence containing it is the object to which it refers, under a mode of presentation. Some have supposed the sense of a name to be a description or cluster of descriptions that singles out its referent.[1] If that's so, there is no pressure to deny that there can be sense without reference: a name would remain meaningful when empty in virtue of its descriptive content.[2] Evans follows Michael Dummett, however, in seeking to dissociate descriptivism and Fregeanism.[3] The way that Evans accomplishes this is by claiming that singular terms are *Russellian*, that is, meaningful only if they refer.[4] This claim can be squared with a commitment to senses when these are construed as object-dependent: as a way that two or more subjects need to be *thinking of the referent* in order for communication to succeed.

Underlying this suggestion is a more foundational idea that Evans adopts from Russell, namely that there are two importantly different kinds of thought or talk about a thing: *by description* and *by acquaintance*. For Russell, acquaintance is a distinguished relation that a subject can bear to her own sense-data, universals logical forms; for a time he also counted the self among the objects of acquaintance.[5] Singular thoughts—the kind usually expressed using singular terms, namely demonstratives, indexicals, and proper names—are acquaintance-based, which is to say that the relevant way of thinking of a particular is founded on the subject's relation to it rather than on her grasp of a description that happens to single it out. But one cannot be acquainted with what doesn't

---

[1] Among them Strawson (1959), Ayer (1964), Searle (1969), and Kripke (1972).

[2] As Frege at times seems to allow, e.g. in the context of his aforementioned discussion of "mock thoughts" in (Frege 1979).

[3] See for instance (Dummett 1973, 97–8).

[4] With the exception of descriptive names, discussed in chapter 1 of *Varieties*. For convenience I shall use the label "singular term" as if it excluded these expressions.

[5] On the objects of acquaintance, see for example (Russell 1905, 873; 1910–11, 108). Russell rejects acquaintance with the self in his (1913); he appears to give up on acquaintance altogether while imprisoned for his activism, in July 1918 ([1918] 1986, 255).

exist; and so one cannot have a genuinely singular thought about what doesn't exist, either.[6]

One of the ways that Evans deviates from Russell is by insisting that the objects of perceptual acquaintance are not sense-data, but ordinary things. Extending the reach of acquaintance in this way, however, allows for a possibility that Russell's picture does not: the possibility of illusions of singular thought.[7] A subject can falsely take herself to be acquainted with an object, and attempt singular thoughts about it on that basis, with the result that, unbeknownst to her, she fails to have a thought at all.[8]

To clarify the connection between this last idea and the claim that singular terms are Russellian, note that a *thought*, on Evans' usage, is a propositional attitude; a singular thought, in turn, is an attitude whose content is singular. Since the senses of singular terms are object-dependent, a sentence containing one or more singular terms that don't refer fails to express a proposition.[9] There is then no content for a subject to grasp or entertain, and in this way the way the defect at the sentential level ramifies into the mind: no object, no proposition, no thought.

Let's pause to situate this view against the backdrop of Prior's challenge. It falls under the heading of the disjunctive strategy in virtue of giving up *(c)*: we have two radically different accounts of what is going on when there is an object of thought and when it isn't. But we also need to draw on our reframing of the challenge's claims to state the view correctly. It's not strictly accurate to characterize Evans' position in the original terms of *(a)* and *(b)*: that X's thinking of Y is a relation to Y when Y exists, but X's thinking of Y is not a relation to Y when Y does not exist. On his view, the empty case is not even a case of *thinking of something*. The good case is a relation, the empty case is just failing to have a thought. This is a possibility that we made room for by retreating, in §2.1.1, to the observation that one *says* "X thinks of Y" when Y refers and when it doesn't.

Note that the claim in Prior's challenge and the Evans-McDowell characterization of the empty case are both stronger than the disjunctive strategy requires. A purely

---

[6] Prior to about 1906, Russell endorsed a *dual relation* theory of judgment, as relating a subject to a proposition. He abandons this model in 1910 in favour of a *multiple relation* theory of judgment. Acquaintance does not figure in the latter account as the basis of a subject's grasp of a singular proposition; the view doesn't trade in propositions at all. Instead, judgment is a multi-term relation between a subject and various properties, objects, and relations with which she is acquainted (which would have figured as components of propositions on the dual relation account). These are grouped together in an act of understanding, resulting in a unit that can be assessed for truth or falsehood. See (Russell 1906; 1910; 1913), as well as the introduction to *Principia Mathematica* (Whitehead and Russell 1910, 44).

[7] Recall the discussion of §2.1.2.1 above and see (McDowell 1987).

[8] (Evans 1982, 46), (McDowell 1977, 173; 1982b, 304).

[9] It may help here to contrast Evans's view with an account on which there can be sense without reference. On the latter understanding, the sense of an expression ensures that the expression is meaningful even if it doesn't refer; for the same reason, it can still make a contribution to the sense of the sentence in which it appears. The result is a sentence that (as we would now say) expresses a proposition. Evans maintains that the sense of a referring term is a way of thinking *of the referent*, where this requires that there be a referent to think about. Since the sense of a sentence is a function of the senses of its parts, a sentence one of whose parts lacks both reference and sense is radically defective: it doesn't manage to express a thought at all.

negative characterization of what happens when the object of thought does not exist would be simply that this is *not a relation*, that it's *not a singular thought*. This is consistent *both* with the subject's having a thought that is not singular (what Prior's way of putting things assumes) *and* with the subject's not having a thought at all (as Evans and McDowell would claim). Here I take the Evans-McDowell account of the empty case on board for expository purposes, with the intention of vindicating it progressively. It's also a heuristically useful starting point, in that working with this view will make it obvious that the account of the empty case I pull together from Evans's writings in no way relies on the attribution of a thought, singular or otherwise, to the subject—which is one of the view's most striking features.

3.2 Three challenges

By all appearances, empty singular terms *are* meaningful, apt to figure in communication and understanding and to do so whether one is informed or ignorant of their status as empty. By all appearances, one who attempts a singular thought in the absence of a suitable object winds up with a thought, a thought internally just like the thought one gets in a successful episode of singular thinking. Hearing a child say, "I am worried that Santa Claus has forgotten my gifts," we gain insight into her state of mind and a rationalization of her actions; yet we are supposed to deem this utterance fundamentally defective and deny that there is a worry that the child expresses by uttering it.

We can extract from these considerations three closely related challenges to the position just introduced, versions of objections that might be levelled at the disjunctive strategy more generally:

> *Concerning the subjective perspective:* If a failed attempt at singular thought produces no thought at all, why does this fact escape the notice of the subject?
>
> *Concerning language:* Why do empty singular terms and the sentences containing them strike us as meaningful?
>
> *Concerning psychological explanation:* What explains the subject's behaviour in the empty case, and why are we so willing to ascribe beliefs and other propositional attitudes if she has none?[10]

A first line of defence is to question whether we need to say more than that the subject cannot discriminate her situation from that of entertaining a singular thought. (As posed here, the challenge about the subjective perspective seeks after a positive feature that explains the subject's ignorance, whereas the negative disjunctivist would deny that invoking such a feature is necessary.) The response that I am concerned with in this chapter, however, is an account of the empty case that independently addresses all three of the queries just listed.

---

[10] A line of criticism pursued by Carruthers (1987), Segal (1989), and Noonan (1991), with McDowell as an intermediate target; see §3.6 below.

As Evans realizes, the cases to be addressed fall into two main kinds. In the first, a subject uses an empty term without recognizing that it is empty, and suffers an illusion of singular thought. Such is the case of the child who wonders when Santa Claus will bring the gifts, or the wanderer in a desert who, taken in by a mirage, asks: "How long until we get to that oasis?" (Let's call these subjects "innocent."). The second kind of case involves what Quine (1960) called the "conniving" use of empty terms, their deliberate use by speakers who *know* them to be empty. How can someone who realizes that the name "Santa Claus" is empty take up the point of view of, communicate with, and understand the claims of a subject who doesn't realize this? It's these cases that Evans is under the most urgent pressure to explain in a way that accords with his antecedent commitments. To see why, we need some background concerning what Evans calls "informational states," and their relationship to belief, singular thought, and understanding.

3.3 Thought and information

We gather information about the world through perception, transmit it through testimony, and retain it in memory. Perception, communication, and memory together form an *informational system* that undergirds our cognition (122). The basic constituents of this system, per Evans, are *informational states*. These states are contentful representations, though they may be non-conceptual, as Evans thinks is true of the states generated when we perceive.[11] Crucially, informational states are *not* propositional attitudes. Instead, they have a role much more like that commonly attributed to mental files.[12] For instance, Evans considers the bundling together of information as concerning a single object to be the "central feature" of the information system (126), and which object, if any, an informational state is *of* is determined causally and not satisfactionally (i.e., *not* by reference to the properties that the informational state attributes).

In Evans' framework it's *thoughts* that are propositional attitudes; and he conceives of these in the way that Russell did early on, namely as a dual relation between a subject and a proposition.[13] For *S* to believe that *p* is then for *S* to stand in the believing relation to the proposition *p*, and so on for other attitude-types. One reason this distinction matters is that it clarifies the no-thought claim. Having a thought is a matter of standing in a relation to a proposition, so the absence of a (singular) proposition is enough to rule it out that the subject has a (singular) thought. More importantly, by taking notice of what distinguishes thoughts—propositional attitudes—and informational states, we discover a way that Evans is able to be concessive.

The representational significance of an informational state is a consequence of two main features, its *content* and its causal *source*. Evans illustrates these features by an

---

[11] In contrast, because testimony is linguistic, testimonial information states are conceptual (Evans 1982, 123n3). Further discussion in §5.2, §6.3 and §7.4 of the same text. Presumably the states preserved in memory are of both kinds.

[12] Confusing matters somewhat is the fact that singular thought has sometimes been identified with thinking via a mental file, including by Recanati (2012), who intends this as a proposal in an Evansian spirit.

[13] See n9 above.

analogy with photographs. The informational content of a photograph might be specified using an open sentence with a variable for each object in the photo, for instance

Red(x) & Ball(x) & Yellow(y) & Square(y) & On Top Of(x, y)

to represent the content of a photo of a red ball on top of a yellow square (125). The object that serves as input to the camera when a photo is produced is, in turn, what the photo is *of*. That source object is also what we would need to compare the photograph with in order to assess the camera's accuracy.[14]

Evans claims that the social informational system constituted by several communicating subjects has this same structure: it "reliably [produces] states with a content which includes a certain predicative component, or not, according to the state of some object" (125). More concretely: exercises of our perceptual capacities and the transmission of information through testimony reliably generate states that do or don't predicate, e.g., redness of an object according to whether that object is red. Our informational states are by and large responsive to how the world is.

A key contrast is between informational states whose content is *purely existential* and states whose content is *singular*. Both can incorporate information deriving from a particular object, and be true or false according to how things stand with that object. Evans denies, however, that states with purely existential content can count as being *about* the object that their information derives from.

Consider first the existential case. I see a bear, x, with one ear missing, and I tell you as much. Hearing this, you gather the information that there is a bear with an ear missing. This content is existential, but still counts as expressing information from x, the bear I saw. The reason is that to assess the predicative aspect of the claim—the part about the missing ear applies—we would need to check in with x. It's likewise with x that the content must be compared to decide whether you, having formed beliefs on the basis of the information I conveyed, now possess knowledge. Your resulting beliefs about x embody information from x. Crucially, however, your beliefs are not thereby *about* x.[15] You are not in a position to have singular thoughts about the object from which your information derives (128).

Getting to singular thought requires something further, as an example from Evans nicely illustrates. The setup is as follows:

> *S* and *A* were in the habit of going hunting together in their youth. On one of their hunting trips, they saw a dazzlingly beautiful bird perched in a pine tree. Years later, *S* (the speaker) may advert to this incident, and say something like: 'Do you remember that bird we saw years ago? I wonder whether it was shot. (307-8)

---

[14] I follow Evans's line here for expository purposes, but it should be noted that he is exercising some creative license in developing his analogy. First, whereas we do assess drawings and paintings for their faithfulness to their subject-matter, we generally think of photographs as just capturing the appearances presented to the camera. Those appearances may be deceptive, but that's not a fault of the photograph *per se*. Second, Evans asserts but does not establish by argument that the content of a photograph should be specified without reference to the individuals that it depicts (1982, 125n10).

[15] Contrary to the "photograph model" of aboutness that Evans rejects in chapter 3 (1982, 76).

The demonstrative ("that bird") and subsequent anaphora ("it"), as well as the talk of *remembering*, all signal that we are in the territory of information-based, i.e. singular thought. Without yet recalling the incident, however, *A* may go along with the story:

> *S* may . . . indicate a very similar bird, perhaps in a photograph, or in the wild. *A*, knowing *S* to be trustworthy, will believe all that he says, and will of course be prepare to bring his belief (which may become quite rich and detailed) to bear upon *S*'s original remark. (308)

But the belief that there is a bird which is *F*, *G*, etc. is not enough. *A* only understands what *S* has said, Evans insists, when *A* "*remembers* the bird", that is, retrieves the *"right* information" (*ibid.*, emphasis original). When there is no object, there can be no information derived from it, and so no possibility of understanding statements of the information-invoking type, should they occur.

By "information," however, Evans also means *mis*information.[16] States of the informational system can purport to represent an object without actually tracing back to that object, or to anything at all (in Evans' terms, they may be *about* an object without being *of* it). Information can issue from false testimony, from hallucination; it might be a bit of fiction misclassified as autobiographical memory. These states, moreover, are "belief-independent" (124): an illusion persists though one recognizes it as such, one's "being placed in the appropriate informational state by someone telling us a story does not depend on our believing the story to be true" (123), and so on. Informational states, then, do not depend on the existence of or one's actual acquaintance with the object they are about, nor on one's believing that there is any such thing.

We've already seen that Evans considers understanding (the grasping of a proposition), demonstrative thought, and the information-invoking use of singular terms to be object-dependent. Now we can add that they contrast with informational states in being *belief-dependent* as well: they require a conviction, on the part of the subject, that the relevant particular or particulars exist. This clarifies what goes wrong in the two kinds of case introduced above. In one—the case of unwitting error—the belief is present, but the object is missing. In another—the conniving case—the appearance of singular thought, understanding, and so on is doubly confounding, because the requisite belief is absent as well.

But here is the promised concession. The fact that informational states are both belief- and object-independent means that *they* can be called on to explain what is happening in both versions of the empty case. Suppose, for instance, that you are hallucinating, but you mistake this experience for a veridical perception. You attempt various demonstrative thoughts about an object that you seem to be presented with. These attempts fail, because there isn't actually an object you're tracking. Of course, it nonetheless seems to you that you are entertaining various demonstrative thoughts. As concerns the level of thought, Evans can endorse a version of the disjunctive strategy: what his opponents would group together as a single kind of mental state, he characterizes as dividing into *singular thoughts* and *merely putative singular thoughts*. But Evans is not a disjunctivist about sensory experience. He holds that perceptual informational states are involved in hallucination as well as in veridical illusion.

---

[16] Compare, for instance, Dretske's (1981) use of the term.

Likewise, the possession of information derived from testimony may be a common feature of successful and unsuccessful attempts at singular thought. In either case, there *is* a positive element of the subject's experience that she is being taken in by.

Evans complains that the belief-independence of informational states is what "casts such a shadow over the theory of reference" (327), the reason being that "it makes possible operations which are highly similar to understanding, in the absence of the commitments which understanding involves." (327) But while these operations are a source of resistance to Evans' position, the fact that informational states are belief-independent proves to be the key for making sense of the conniving use of empty singular terms.

3.4 A role for pretence

The resources we need are in chapter 10 of *The Varieties of Reference*. The advertised aim of this chapter is to introduce a novel treatment of negative singular existential statements, such as "Pegasus does not exist."[17] In fact, however, what Evans says about negative existentials is just an application of a more general strategy for dealing with empty singular terms, which in turn has implications that go well beyond the difficulties about language to which it is most explicitly addressed.[18]

The strategy involves pretence. The first person to make some connection between that phenomenon and the use of singular terms that lack a referent is Frege:

> Assertions in fiction are not to be taken seriously: they are only mock assertions. Even the thoughts are not to be taken seriously as in the sciences: they are only mock thoughts. (1979, 130)

For Frege's this is a passing observation about an oddity of natural language, and the idea is not one that he pursues extensively ("The logician does not have to bother with mock thoughts." (*ibid.*)) He is also, in this passage, targeting a narrow range of utterances. The possibility that there might be an issue of general interest here, however, is more apparent in the following remark:

> I am convinced that the idea I associate with the worlds 'my brother' corresponds to something that is not my idea and about which I can say something. But may I not be making a mistake about this? Such mistakes do happen. We then, against our will, lapse into fiction. (Frege 1956, 306)

Though the notion of an unnoticed and unwanted lapse into fiction has a strange ring, the passage is suggestive nonetheless. Evans charges Frege, in his "use of the notion of fiction," with "wrongly [directing] our attention to just one case in which it is as if a singular term refers to something" (Evans 1982, 30). That may be so, but we can also find in these remarks the germ of an idea that Evans will exploit to great effect, namely that the things we say and think when we are self-consciously fabricating and the errors

---

[17] Dummett (1996) finds a problem with the letter of Evans' account, which stems from its reliance on counterfactuals. See Walton (1990) for another way of spelling out the view, which instead exploits the notion of a "kind of referring attempt" (427–9).

[18] Wiggins (1995) hints at this broader significance.

that we fall into in our efforts to engage with the world might fruitfully be considered, and to some extent treated, together.

The solution that Evans develops to his problem about the conniving use of empty singular terms is inspired specifically by Kendall Walton's early work on the role of make-believe in our engagement with the representational arts.[19] Children engaged in pretend play, as Walton observes, regularly make assertions that they know to be false (1973, 287). For example, in a game of mud pies: "There are three pies here," "That idiot man stepped on one of our pies" (287). Despite this proliferation of falsehoods, make-believe is not a free-for-all: some contributions strike participants as acceptable and are treated as such, others not. Barring an allowance for invisible pies, it would be *wrong* to say, looking at a table on which are located only two globs of mud, "There are three pies here." As with ordinary assertions, a demand for justification might be raised, and providing that justification means working within or modifying the boundaries of the game.

Make-believe is distinguished from other employments of the imagination by the involvement of *props* (1990, 12). A toy stethoscope that a child uses in pretending to be a doctor is a prop, but so is a mud pie, or a stump in a game of make-believe in which it is pretended that stumps are bears (21); what serves as a prop need not have been created for that purpose, nor need it be an artefact at all. What is characteristic of props is rather that they "*generate fictional truths*" (*ibid.*; emphasis original). They do so in combination with what Walton terms "principles of generation" (38), e.g., the principle that

> (1) Where there is a glob of mud *there is a pie.*

where the asterisks indicate truth-in-the-make-believe.[20] If (1) is in effect, then the presence of a mud pie at location *l* makes it the case—generates the fictional truth—that *there is a mud pie at *l**. This mechanism has the interesting consequence that participants in a game of make-believe can be ignorant of or mistaken about truths in the game; they can *discover* states of affairs that they have not themselves stipulated.

Props also serve, in make-believe between multiple subjects, a coordinating purpose (1990, 19). In "collective imaginings" that do not involve props, correspondence among the imaginings of the participants is secured in an active and deliberate fashion, with stipulations and agreements to imagine this way or that. The result is not "mere correspondence in what is imagined," but a sort of "common knowledge" correspondence:

> Not only do the various participants imagine many of the same things; each of them realizes that the others are imagining what he is, and each realizes that the others realize this. Moreover, steps are taken to see that the correspondence obtains. And each participant has reasonable expectations and can make justified predictions about what others will imagine, given certain turns of events. (18)

Props can play the same role insofar as they strike people in characteristic ways. There is a shape that a cloud can have, for example, such that it will prompt many or most

---

[19] (Walton 1973; 1978). See also the discussion of Evans's ideas in (Walton 1990, chap. 10).
[20] Walton has since abandoned this notation, which I persist in using for ease of exposition.

perceivers of it to imagine a running horse, and to expect that others would imagine the same.

As Walton emphasizes, not all pretending takes the form of overt behaviour. It may be "fictional not just that our eyes are aimed in certain directions [...] but also that we notice things, that we make discoveries and acquire beliefs." (240) Linguistic pretence also has a psychological component, as will become important when we return to Evans:

> Verbal participation is not just a matter of fictionally mouthing certain words. Fictionally we tell people things, expecting them to believe what we say, or at least to understand. It is fictional, ordinarily, that our various verbal and visual actions involve intentions, expectations, desires, hopes, and beliefs of the kinds that attend similar actual actions. (*ibid.*)

On Walton's account, psychological states can themselves serve as props, enriching the make-believe and combining with principles of generation to yield fictional truths.

How this would figure in a simple game of mud-pies may not be immediately clear, but Walton's view is that our adult appreciation of works of representational art—novels, paintings, films, certain pieces of music, and so on—is a development of the same activity. One much-discussed aspect of our engagement with fiction is the emotional responses that this can apparently engender, and it's here that Walton finds a special role for psychological props: they can make it fictional that one feels a certain emotion.

Suppose that an actor is in a scene in which his character prepares to perform some terrifying feat, and that the actor, incidentally, suffers from stage fright. The actor may then use his own jitters as a prop as he imaginatively takes on the character's state of mind. As the actor becomes aware of an increase in heart rate, fictionally the character becomes aware of an increase in heart rate; as his thoughts circle anxiously, fictionally the character's thoughts circle anxiously; and so on.

This last idea comes to the fore in Walton's "Fearing Fictions," through a discussion of the following case:

> Charles is watching a horror movie about a terrible green slime. He cringes in his seat as the slime oozes slowly but relentlessly over the earth destroying everything in its path. Soon a greasy head emerges from the undulating mass, and two beady eyes roll around, finally fixing on the camera. The slime, picking up speed, oozes on a new course straight toward the viewers. Charles emits a shriek and clutches desperately at his chair. Afterwards, still shaken, Charles confesses that he was "terrified" of the slime. (1978: 5)

Charles's "muscles are tensed, he clutches his chair, his pulse quickens, his adrenalin flows" (6). He is in a physiological condition that produces sensations and feelings that we associate with fear. But is he really afraid? Walton answers "no." Charles does not *believe* in the slime. He does not *behave* as though he were under immediate threat (calling emergency services, running away, etc.). He is entertaining himself by engaging in a game of make-believe that takes advantage of the responses of his autonomic nervous system and the impact that these have, whatever he really believes about his

environment, on his conscious experience (13). That activity puts him in a distinctive mental state, which Walton labels "quasi-fear" (6).[21]

Critical reaction to Walton's paper has focused on the claim that Charles does not feel real fear, but that is not the primary moral of the example.[22] Walton provides an account of why we describe Charles as afraid, and are in fact licensed in doing so, that doesn't *presuppose* the reality of Charles' emotion. The obvious explanation of our linguistic behaviour would be that we are identifying an instance of fear and labelling it as such. Walton's alternative is that we call Charles' emotion "fear" within a certain conversational frame, one established by a tacit game of make-believe. When we say things like this, we are ourselves engaged in the pretence Charles initiates around the horror movie. Relative to the actual state of things, we may speak falsely; but in the scope of the make-believe, what we say is true.

When he comes to consider another kind of case, Walton goes further. Noting that we often describe fictional narratives by reporting them as though they are factual, he makes the following suggestion:

> When Sally says, "Tom Sawyer attended his own funeral" in pretense, it is *fictional* in her game that there is someone whom she calls Tom Sawyer and who she claims attended his own funeral. We need not suppose that there really is any such person. We need not suppose even that there is someone, or some character, about whom Sally *pretends* to speak, nor that there is a proposition she intends to assert. Her utterance is fully intelligible as an act of participation if we take it to be *merely* fictional, not true, that her words express a proposition, and that they express a proposition about someone whom she calls Tom Sawyer, and that she asserts such a proposition to be true. (396)

What's striking in this passage is the suggestion that Sally's utterance of "Tom Sawyer attended his own funeral," though perfectly apt, expresses no proposition. So it's not only literal falsehoods that can be redeemed by their application in a pretence (e.g., "this is a cat," demonstrating a crumpled sweater), but sentences that are strictly meaningless as well.

This commitment sharpens the contrast between Walton's approach to fiction and another that superficially resembles it. David Lewis (1978; 1986) calls on possible worlds to provide the propositions that seem to be, and on his view really *are*, expressed in fictional

---

[21] This, at least, is what Walton sometimes seems to intend the label to pick out; at other moments he uses "quasi-fear" to talk about the psychological and physiological common-ground between a case of fearing and a case of merely pretending to be afraid. I'll adhere to the first usage for continuity with the way Evans takes up this aspect of Walton's view. Walton has also emphasized to me in conversation that he has moved away from using the term "quasi-fear" in the intervening years because of the confusions that it has engendered.

[22] The question of whether our responses to fiction constitute genuine emotions has received ample discussion, with most insisting, unlike Walton, that Charles is genuinely afraid. See for instance (Lamarque 1981), (Stocker 1987), (Greenspan 1980), (Carroll 1990), (Neill 1991), (Moran 1994), (Matravers 1998), and (Prinz 2004). (Friend 2017) provides an overview of the debate.

works and the associated discourse.[23] Tom Sawyer isn't part of the actual world, but he, and the states of affairs of which he forms a part, nonetheless exist (in some possible world) to be referred to and described. On this view, statements like Sally's feature a hidden operator ranging over possibilia, "in the fiction" or something to that effect, thanks to which they get to count as literally true.

Walton proposes to understand Sally's utterance in a fundamentally different way. An operator like "in the fiction" needs entities to operate on, and he is concerned precisely to avoid a commitment to fictional objects, however construed.[24] Instead, the account is pragmatic in the following sense: statements that would be defective if taken literally count as legitimate and cooperative contributions because we are not *using* them to make, or interpreting them as, ordinary assertions. Indeed, we are not making assertions at all, but merely pretending to do so. While these uses of language look quite ordinary (compare "Barack Obama attended the funeral of John Lewis"), they belong to a distinct activity and must be interpreted in that light.[25]

3.5 The shape of the solution

The claims of the preceding section come together with the outlines of Walton's theory of make-believe to resolve the various questions that Evans faces as a result of his distinctive commitments:

> a proper appreciation of the nature and extent of linguistic pretence [...] enables us to hold on to the insight that singular terms are (generally speaking) Russellian, while taking a realistic and credible view of phenomena apparently inconsistent with that insight. (340)

Evans finds in Walton's theory a way of explaining the appearance of singular thought, singular reference, communication, and understanding without conceding that any of these are actually present in the empty case. The outcome is not just that we can consistently maintain that singular terms are Russellian. It's that we can hold onto a conception of certain mental states as irreducibly relational, as dependent on the things in the world that they concern; we can stave off the outcome that the skeptic tries to drive us towards, which places mind-independent reality beyond our cognitive and our experiential reach.

To bring out what lies a bit submerged in Evans' telling, it's helpful to reconstruct his position with the case of Charles and the slime in mind. Charles has sweaty palms, a

---

[23] For discussion of the relative merits of the make-believe and possible-worlds approach, see e.g. (Currie 1990) and (Kroon 1994). Compare also (Priest 2005) and (Berto 2008; 2011), both of which invoke impossible as well as possible worlds.

[24] See the first pages (Walton 1990, chap. 10 sec. 1). Note, however, the possibility of differentiating two uses of the names of fictional characters: one that we would classify as the use of an empty name (e.g., "Sherlock Homes lived on Baker Street") and another that we could construe as referring to an abstract artefact, the product of an author ("Doyle created Sherlock Holmes before a business lunch").

[25] The idea that some sentences containing empty singular terms express no proposition except as considered *internal to a pretence* is one that Kripke (2013) introduces, and makes several applications of, in his 1973 Locke Lectures. Evans and Walton are the focus of my discussion because the line I want to take on these issues dovetails more closely with their accounts.

fast-beating heart, a feeling of agitation. He may say "That slime is going to get me!" and "I'm terrified!" He does not feel real fear, or so Walton suggests. Evans, however, is not specially interested in emotional responses to fiction; the mental state that troubles him is *understanding*. He wants to argue that what looks like understanding in the empty case is instead, on the model of quasi-fear, a state of quasi-understanding (1982, 363).

We know that responses like an elevated heart rate and sweaty palms arise automatically in response to stimuli; in Evans' terms, they are belief-independent. And it's these responses that serve as props in Charles' make-believe about the slime, helping to make it the case that he is quasi-afraid—fear itself being a belief-*dependent* state, according to Walton. The parallel with the cases that interest Evans should now be coming into view: in a similar way, our belief-independent informational states might be props that help make it fictionally the case that we have singular thoughts and achieve understanding. Or, to describe the situation from outside the pretence: that make it the case that we quasi-believe and quasi-understand.

Not just any act that might be called "pretending to understand" can play the role required here. One doesn't achieve special insight by nodding along vigorously to a lecture on quantum physics that exceeds one's comprehension.[26] But the proposal at issue here is that one recreates, within a pretence, the kinds of conditions that make understanding possible—"internal" as well as "external" conditions. The result is not (merely) the imitation of certain behaviours characteristic of understanding, but an analogous operation of the mind. It also bears emphasizing that we have been working with a narrow notion of "understanding," as the grasping of a proposition. To say that the pretence issues in quasi-understanding is thus not to deny that the exercise also has straightforward or "real" (i.e., extra-pretence) epistemic dividends.

One of the crucial ideas here is that not only our own inner states, but the inner states of others, could serve as props in the way that we saw Charles's awareness of his own racing heart serve as a prop. *Your* hallucination and *your* false beliefs might make it fictionally the case, even in a game of make-believe in which I am the sole participant, that you are having demonstrative thoughts about an object you are seeing. Whether the informational states are yours or my own, pretending in this way comes easily, for having resolved to pretend that the (mis)information is accurate, "one can let the automatic and habitual responses of one's cognitive system take over and produce the make-believe thoughts, emotions, and reactions which playing the game normally requires" (359).

Recall that understanding is ruled out on two grounds when one knows that a term is empty: first, because there is no proposition to grasp, and second, because understanding is belief-dependent.[27] But pretend there is an object that your information derives from, and within the make-believe there will also be singular propositions concerning it. Pretend there is an object you are acquainted with, and you will tend to be pretending also that you believe it exists; the conditions are then met, within the pretence, for information-based

---

[26] When Mark Sainsbury describes Evans as appealing to an *imitation* of understanding, he seems to have cases like this in mind (2000, 65). If that was Evans's proposal, Sainsbury's concerns about it would be totally just.

[27] This is a place where quasi-fear (as described by Walton) and quasi-understanding come apart. Walton is happy to grant that we can be afraid of what does not exist, as long as we falsely believe it does. Evans puts both an *existence* and a *belief* requirement on understanding.

thought, communication, and understanding. Informational states, used as props, scaffold the non-believer's pretence.

Informational states are distinctly suited for this role because of their existential import. A hallucination represents *that there is a noise coming from the ceiling*. A bit of misleading testimony represents *that there is a man named "Santa Claus" living at the North Pole*. Simply by treating these states as veridical, one becomes engaged in what Evans calls "existentially creative" make-believe (*ibid.*): one doesn't merely ascribe novel properties to things in one's environment (as in the game of mud-pies) but pretends that the world includes objects that aren't really there.

That no separate stipulation (to the effect that such-and-such a thing exists) is required is a feature of the proposal that needs to be underlined. Singular thought involves the uptake of information emanating from an object, which is thereby determined as what the attitude is about. Consider an ordinary case of perceptual demonstrative thought, which the speaker expresses by saying, e.g., "That thing is ugly." To grasp the speaker's meaning, the hearer must also be in a position to perceptually identify the object in question. For our account of the empty case to be plausible, we need a story about how pretence makes an analogue of this kind of coordination—thinking of the same thing *and* in the same way—possible.[28] A pretence initiated with a description that one imagines the would-be object of the would-be singular thought to satisfy will not clear this bar, because distinct things can be qualitatively identical. A description that exploits ways that the subject is (fictionally) embedded in her environment, e.g., "whatever is the source of her information," makes the target determinate, but falls short as concerns the requirement to think of the object in the same way: entertaining a description of this kind is not normally involved in thinking (having singular thoughts) about the things one is acquainted with.[29] To mirror the ordinary case, in which we get onto a single object under the same mode of presentation, the pretence must be grounded in the subject's misinformation. Treating *that* information as accurate lets it determine how the world fictionally *is*—taking advantage of the special feature of props I highlighted in (§3.4).

Notice that there is a close parallel between the situation of the conniving and the unwitting users of empty terms, or "thinkers" of (non-existent) singular thoughts, on Evans' picture. The cognition and behaviour of the unwitting subject, too, are scaffolded by informational states. The central difference is that she winds up in a state of quasi-understanding because she is misled, whereas the conniving subject does so in virtue of engaging in pretence. The unwitting subject automatically treats her informational states in the usual fashion, whereas the conniving subject resolves to do so. That the conniving and the unwitting user of an empty term can coordinate in this way contributes to explaining why "not all uses of empty singular terms need be utterly *opaque* to an undeluded and suitably informed observer" (340).

3.6 Revisiting the challenges

---

[28] More about the relationship among attitudes in chapter 4 (§4.2).

[29] "In giving the content of [a singular] thought, we need to specify the object in a way which mirrors the mode of identification which the subject employs; hence, *not* as the object which bears the appropriate causal relations to the information on which the thought is based." (Evans 1982, 139)

With these materials on the table, we are now in a position to formulate responses to the three questions I introduced in §3.2:

> *Concerning the subject's perspective:* If a failed attempt at singular thought produces no thought at all, why does this fact escape the notice of the subject?
>
> *Concerning language:* Why do empty singular terms and the sentences containing them strike us as meaningful?
>
> *Concerning psychological explanation:* What explains the subject's behaviour in the empty case, and why are we apparently willing to ascribe beliefs and other propositional attitudes?

As above, the immediate answer to the first question is just to allow that subjects can be wrong about the nature of their own mental states. But we can say more about how such a confusion might arise. The subject *does* have information states in both cases, which can represent objects and situations that don't exist, which are bundled together according to whether they are presumed to emanate from the same thing, and which she will treat as relevant to the thoughts she (unsuccessfully) attempts. In addition, we have an account of why it's so easy for outsiders imaginatively to take up this confused perspective, though recognizing it as such; and why, in imagining it, we don't *think* of it as gappy either. We imagine the subject's perspective by imagining that the world is as she takes it to be, and in such a world, the subject *does* have a singular thought.

This last point is crucial for counteracting charges of phenomenological implausibility. The claim "no object, no thought" concerns what a thinker is in fact, objectively speaking, able to do. It is not a claim about what it is like for the subject in the empty case. Of course, what the thinker is able to do has, in this instance, consequences for their mental state, and so it can be hard to see how it could fail to have implications for her conscious life. But we are allowing that there is a distinction between appearance and reality *within* the mind. Once that move has been made, the question of what mental state someone is in and the question of what it is like to be them will sometimes receive divergent answers. Accurately characterizing the subject's point of view is its own task, and it's the demands of that task in particular that pretence rises to meet.

As concerns the second challenge: empty singular terms strike us as meaningful because we have a way of accommodating the lack of a referent. We pretend that there *is* a referent, and do so simply by taking some (mis)information at face value. Both Walton and Evans would insist that you do this sort of thing automatically when you read that once upon a time there was a dragon who guarded a trove of stolen valuables… etc., at least if you are invested in enjoying the story. One can explain in the same way cases that have nothing to do with works of fiction. Consider, for example, the experience of listening with a mix of empathy and suspicion as a friend describes a perfect Tinder match who keeps asking for money but refuses to meet in real life. However much you may doubt the existence of the person being described to you, you can recruit your capacities for thinking about individuals (accumulating a list of properties associated with an object, tracking it

through changes, etc.) by engaging in pretence—where again this is a matter of treating their misinformation as true for the purposes of the exchange.

Suppose that Caroline suffers a hallucination as of a crabapple tree, but does not realize that she is hallucinating. As far as she can tell, she is forming various perceptual demonstrative (and hence singular) thoughts. We expect Caroline to act in certain ways: to say things like, "that's a nice tree," to reach out to pick a fruit, and so on. Unless and until she realizes that she is hallucinating, that is, we expect Caroline to act in just the same way as someone who is seeing a tree and having singular thoughts about it. But what are we to cite in the explanation of her actions? The attitudes we would appeal to if they were veridically perceiving are, on the view we're considering, simply absent in a case like this.

Consider, by way of comparison, McDowell's (1977) response to the problem about the explanation of behaviour in the empty case. His idea is that we can appeal to the subject's higher-order belief in lieu of the singular thought we would otherwise appeal to:

> the belief which makes the behaviour intelligible is a (false) second-order belief to the effect that the subject has, and is expressing, a first-order belief correctly describable in the transparent style. This second-order belief is manifested by the subject's action, not expressed by his words. No belief is expressed by his words: they purport to express a belief which could be described in the transparent style, but since no appropriate belief could be thus described, there is no such belief as the belief which they purport to express. (173)[30]

This solution has an unfortunate structure. It works by appealing to an element that is supposed to be both (i) common to the good and bad cases, and (ii) sufficient to explain everything that we want to explain about the subject's behaviour in the bad case. As a result, it's vulnerable to a style of objection that Noonan (1986), Carruthers (1987), and Segal (1989) have variously pressed, which takes the form of a dilemma: either the second-order belief actually plays no role in psychological explanation, or it plays the same role across the good and the bad cases, making the singular thought that was supposed to distinguish the good case explanatorily redundant.[31]

Obviously neither fork is desirable. The dilemma can be avoided, however, by deploying the resources I have been foregrounding in this chapter. Informational states are common to the good and bad cases, but do not meet condition (ii) above. The state of mind of a subject who makes a failed attempt at singular thought is *ineffable* in the sense that there is no suitable belief to appeal to in the explanation of their actions, no literal truth we

---

[30] The topic of second-order mental states might inspire a question about the precise contours of an anti-internalist view of singular thought as relational and object-dependent. Some of our thoughts have other mental states as objects. Are these thoughts dependent on the mental state they concern, in the way that thoughts about things in the environment are dependent on those things? Where we go on this question depends on whether we take the rejection of internalism or the assertion of a dependence between thought and its object to reflect the true essence of the proposal. I am inclined to emphasize the second of these, and so to hold that one cannot have a singular thought about a thought unless that second thought exists and is singled out in a non-descriptive manner. This inclination stems from my scepticism about the theoretical value of the internalism/externalism contrast—and in particular a feeling of unease about the brain-or-body boundary typically invoked in distinguishing these views.

[31] Compare again Martin's (2004) discussion of "screening off" worries in the perceptual context.

can state that explains her reasoning and behaving as she does. Yet we have an alternative recourse in the form of a pretence that exploits her informational states.

When we say that Caroline has reached for a fruit because she want to make applesauce from it, for instance, we are *pretending* to ascribe a singular belief to Caroline. As I explained in in the previous sub-section (§3.5), we initiate this pretence by make-believing that her perceptual state is veridical. Internal to the resulting pretence, there is an object that she is seeing, is having singular thoughts about, and is in a position to reach out and grasp. By means of the pretence, we engage in analogues of ordinary communication (with informed third parties as well as the confused subject herself), coordination, and understanding. This proposal is uniquely suited to escape our dilemma. The account exploits an element of the subject's psychology that is present in hallucination as well as in veridical perception, namely a sensory state (a variety, recall, of "information state" (Evans 1982, 122–123). That state is not itself explanatorily sufficient, however: we work our way up to sufficiency by engaging in pretence. Since this activity involves pretend objects and pretend thoughts and pretend propositions, it introduces (in reality) no further element that would screen off singular thoughts from playing the unique explanatory role that we want to accord to them. Hence the two-way contrast I drew in §2.2.3 between this view (at that point merely foreshadowed) and both negative and positive varieties of disjunctivism, which we can now make better sense of. The view posits information states and explains in positive terms how these bring about, in combination with either the subject's misplaced faith in the accuracy of her information (the innocent case) or a pretence to the effect that the information is accurate (the conniving), the appearance of singular thought in the empty case.[32] Yet the account is in another way more closely aligned with negative disjunctivism, in that it does not attribute a mental state in the empty case that is itself sufficient to account for the facts about the subject's point of view, our use of language, and our capacity for psychological explanation that formed the basis of the objections I have just addressed.

3.7 The bigger picture

As we've seen, one of the ways that Evans elaborates Walton's account is in finding a role for pretence in an activity even more commonplace than the consumption of fictional narratives. Children make a radical gain in predictive power and psychological understanding when they acquire the concept of false belief, and we often have no less interest in and no less a need to engage with a person for the fact that they are wrong about how the world is. Make-believe is a means of gaining insight into a class of mental states

---

[32] I emphasize that this appearance may, in a single case, have a different explanation according to whether we are considering the perspective of the subject who fails to have a thought or the perspective of a second person who is engaging with the former. The innocent subject's ignorance of the fact that she has no thought is explained by her sincerely taking her information state to be veridical (which error can be called upon to account for the subjective sameness that a negative disjunctivist would refrain from characterizing positively); other subjects may be in the grip of the same existentially creative misinformation, or they may recognize it for what it is, in which case they are conniving users of the empty singular term at issue. In the latter case, the persistent impression that the innocent subject has a thought is explained by the automaticity with which we imaginatively adopt the perspectives of other people, which in these cases involves information-based, existentially creative pretence.

that each of us is prone to experiencing on occasion, mental states defined by their subjective resemblance to more familiar, veridical ones: the attempt at singular thought grounded in misleading testimony, the hallucination, the false episodic memory.

I want to lay particular emphasis on the manner in which that insight is achieved. The goal is to understand a mental state; the procedure is to pretend that *something external to the mind* exists, that the *world* is other than it is. This claim about our grasp on the minds of others has a first-personal counterpart in Evans' response to the model of introspection as a kind of inner vision:

> …in making a self-ascription of belief, one's eyes are, so to speak, or occasionally literally, directed outward—upon the world. If someone asks me 'Do you think there is going to be a third world war?', I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question 'Will there be a third world war?' I get myself in a position to answer the question whether I believe that *p* by putting into operation whatever procedure I have for answering the question whether *p*. (1982, 225)

To discover how one believes the world to be, one makes an assessment of the world. To continue engaging with someone that has fallen out of step with reality, one concerns oneself with the world-from-their-perspective.

What do we learn from the necessity of this procedure? The lesson is one that the opposing, 'Cartesian' conception utterly obscures, namely that some of what it construes as contents of the mind are just elements of the world on which the conscious subject occupies a point of view. Of course our intentional states do frequently part ways with reality, and for various reasons: error, ignorance, imagination, play. But if our insight into those states and our ways of talking about them demand a form of pretence in which we imaginatively intervene on the world—if there is no *more* direct intervention on the mental state to be characterized—then psychological realities do not have the independence that the internalist grants to them.

Even independently of the specific views expressed in this chapter, one can recognize we rely on statements about the world in our efforts to characterize many kinds of conscious experience. Take a referring term or a proposition, prefix it with a suitable attitude verb, and you have a mental state ascription with implications, not just about what the subject happens to be cognitively related to, but about what it's like to *be* her. In a discussion of *de re* belief, Stalnaker (2009) says that "[t]o think of something as a believer is to think of it as having a point of view on the world, and to ascribe beliefs to it is to attempt to describe that point of view" (243). We accomplish this by "[saying] what the world is like according to them" (*ibid*.).

There is an unusually direct line, here, from a question about what ordinary language users are doing to issues in philosophy of mind. The standard interpretation of sentences like "He is looking for his keys" and "He is looking for Atlantis" has it that they are apt because they express a proposition that is true. On the pretence account, the latter is apt *despite failing to express a proposition*: apt because true in a game whose purpose and whose outcome is a grasp of the relevant mental state that cannot be otherwise secured. The truth of the standard accounts requires that singular thoughts be object-independent; the truth of the Walton-Evans pragmatic alternative does not. The former is compatible with a view of consciousness as fully amenable to direct

description; the latter demands the negation of that same claim. Yet despite these and other consequential differences between the approaches, the linguistic and behavioural appearances are consistent with either of these radically different alternatives.

Chapter 4: Applications

The discussion of chapter 2 cleared the way for the disjunctive response to Prior's challenge, which affirms *(a)* and *(b)* at the expense of *(c)*. The proposal from chapter 3, that some of our ordinary mental state ascriptions belong to a pretence, radically improves the dialectical position of the disjunctive strategy: when the strategy is elaborated in this way, we find that it does not merely accommodate but actually *predicts* observations that are usually held against it. I'll now bring this finding to bear on two debates, both of which bottom out in questions about the nature of singular thought: one about the centrality of acquaintance (§4.1), the other about what it is for multiple such attitudes to qualify as *about the same thing* (§4.2).

4.1 The criterion of singularity

In recent years there has been a renewal of interest in singular thought, as witnessed by the publication of a couple of volumes of collected essays on the topic.[1] The role of acquaintance in delimiting the singular from the non-singular thoughts is a central point of contention in this literature. For example, Robin Jeshion (2010) proposes that what is essential for singular thought is not acquaintance but significance (in a sense to be explained below); John Hawthorne and David Manley (2012), while admitting that there are "certain paradigm cases […] in which causal or epistemic relations do play a significant role" (26), argue that there are no substantive constraints on singular thought and reference; François Recanati (2012) considers acquaintance to be the characteristic, but not strictly an essential, basis of singular thought.

There are numerous aspects to this disagreement. Hawthorne and Manley, for instance, present a whole series of objections to Evans's position on singular thought, targeting some of his assumptions as well as the verdicts that the view yields on various examples (Hawthorne and Manley 2012, chap. 3). I restrict my focus here to a challenge of the latter kind, according to which our intuitions about certain cases push us away from an acquaintance-based conception of singular thought and towards an internalist one.

4.1.1 Acquaintance

Let's first revisit some themes from the previous chapter. A central distinction in Bertrand Russell's theory of meaning is that between logically proper names and descriptions (1905, 1910-11). According to Russell, the latter do not refer but rather introduce conditions on the truth of sentences containing them. A logically proper name, in contrast, is genuinely referring, and is linked up to its referent without the mediation of descriptive content. Because its meaning *is* its referent, no such term can at once fail to refer and be meaningful. The acquaintance relation is what secures the reference of a logically proper name. (Ordinary proper names, according to Russell, are just disguised descriptions.)

To better understand why this matters, consider the case of sensory experience. As the possibility of hallucination makes clear, one can err in one's judgments about what

[1] (Jeshion 2010), (Goodman, Genone and Kroll 2020).

objects are present in the surrounding environment. But there is no question in such cases—or so the early sense-datum theorists assumed—that one is aware of instances of sensible qualities, say loudness or redness (recall the argument from illusion, §2.2.3). Since these properties are ontologically dependent, one expects them to be instantiated by something which is present to the mind even in hallucination or illusion. Sense-data are introduced as filling this role. Thus while there is such a thing as merely seeming to be aware of a table, there is no such thing as merely seeming to be aware of a sense-datum. One's grip on one's own sense-data accordingly has the security that Russell seeks for the referents of logically proper names; the acquaintance relation, rather than a descriptive content, is also what determines *which* sense-data are referred by a demonstrative that one tokens.

As acquaintance figures in Russell's framework, it sets singular thought apart from descriptively-mediated thought in many interrelated ways. For example, the correct use of genuinely referring expressions is restricted to the possible objects of acquaintance; acquaintance-based thoughts about an object *o* are such that their truth depends on the state of *o*, i.e., they have object-involving truth conditions; there are mistakes that we can make about the world that do not arise with regard to objects of the acquaintance relation and hence which cannot infect our singular thoughts. When we move away from Russell's peculiar commitments, the story becomes more complicated, and one might worry that the concept of singular thought loses its unity and perhaps also its utility in this case.[2] At the very least, if we find that there is a use for a notion of acquaintance and a notion of singular thought outside of Russell's picture, it is possible that the two concepts should cease to be intertwined on their new construals.

It is common ground in recent debates about singular thought that if there is a relation that merits the title of "perceptual acquaintance," it's the one that we stand in to ordinary things. So conceived, acquaintance is still epistemically distinguished: perceiving things is a good way to know about them. But subjects are liable to be wrong about whether this liberalized acquaintance relation obtains, and the immediate upshot is that if we want acquaintance to remain a necessary condition on singular thought, we must accept a thesis by now familiar to readers of this dissertation: namely, that there can be illusions of singular thought.

At this point, two reactions are likely. The first is that this outcome is intolerable in itself; the second is that the category of singular thoughts patently transcends the division between the cases where there is acquaintance and the cases where there isn't. I've already addressed at length the issue of its not being transparent to the subject what kind of mental state she is in. But the second of the two reactions just cited—looking at the data to be accommodated and judging that the phenomenon at issue is *manifestly* object-dependent—is one that we are now in a much better position to push back against.

4.1.2 Mental-file frameworks

It will serve us to have determinate alternative to Evans's account in view. To this end, I'll briefly introduce an approach, which centers on the notion of a *mental file,* that has garnered a lot of attention.

---

[2] For example, Sainsbury lists out seven "varieties of singularity" that different accounts of singular thought might make more or less central (2020, 21–2).

The basic idea behind mental file theories is originally due to Grice (1969), Lockwood (1971), Strawson (1974), and Perry (1980), who analogize some of our cognition to the use of file folders for organizing documents. File folders are commonly used to collect information about a topic or (in medical contexts, for instance) about an individual. The folder is identified by its label, say a name of a person whom it concerns. One adds to it whatever new information one collects about that person and takes out what is discovered not to belong (because it's false, or about someone else, or about nobody at all). Files may be merged, as when two of them turn to be about the same person under different names, or split, as in the converse scenario. It seems like we perform similar operations in thought, particularly as we track objects through time and across changes to their properties. Hence the concept of a mental file.

To embrace this description is simply to hold that files are an illuminating metaphor for understanding our thought about particulars. Some developments of the idea transform it, however, into a substantive theory of our cognitive architecture. The central thesis of such a theory is that the mind employs mental analogues of file folders to organize information from objects. Why think this is so? I'll outline a few reasons.

Some mental file theorists take inspiration from empirical evidence for similar mechanisms in perception. Various of the low-level activities the visual system engages in have been explained by appeal to *object files*, which group properties together by attributing them to a single object. These files work together with so-called "visual indexes," or what Pylyshyn (1989) called "fingers of instantiation" (FINSTs), which are posited to explain our tracking of objects over time. These indexes encode spatio-temporal information. The anchoring of an object file to a visual index means that the file can persist—and an object be tracked—even through changes to all of the properties in the file. Object files and mental files are not the same thing, but one might nonetheless be tempted to bolster the mental file idea by pointing out that systems that marshal information in similar ways are operative at lower levels of cognition. Certain theorists also draw a connection between the two types of files. For instance, Jeshion (2010) suggests an association between mental files and FINSTs, claiming that what makes thinking from a mental file singular is the fact that "thought with mental names and demonstratives is ontogenetically rooted in the coupling of them with FINSTs" (135).[3]

There are also features of our thinking about objects—singular thought in particular—that could be explained by mental analogues of file folders. One of these is the pattern of reasoning called *trading on identity*:[4]

> *N* is *F*.
> *N* is *G*.
> So, there is something that is *F* and *G*.

Here one infers from two pieces of information, each predicating a property to an object, that the properties in question are co-instantiated. This move is rational if both pieces of information were in a single file, for to be in a single file is to (be thought to) concern one and the same object. A commitment to identity is thus built into the way the information

[3] See King (2020, 93–4) for a critique of this claim.
[4] Campbell (1987).

is organized and does not need to be made explicit (e.g., be included *among* the information about which one reasons).

Then there is our capacity to track objects as they undergo changes, however dramatic or total those changes may be. If we were managing to think of an object in virtue of its satisfying a description or a cluster of descriptions, such changes would loose our hold on it. Or consider the way that our judgments of quantitative sameness and difference reflect a sensitivity to spatiotemporal continuity, in addition to the way that properties are distributed. For example, breaks in continuity are interpreted differently according to whether or not the interruption appears to be the product of occlusion (the blocking of one object by another, in this case only momentarily).[5] A physical file folder is the sort of thing that exists independently of its contents, having information put into and taken out of it, with the same name on it all the while. Again, the analogy seems promising: these are features well suited to accommodating the phenomena just cited.

Reasoning in a way that takes identity for granted, tracking objects through changes and across space and time—these are features that paradigm cases of singular thought exemplify. The potential of mental files to explain them accordingly suggests a possible association between *having a singular thought* and *thinking by means of a mental file* (a mental files is supposed to be a type of representational vehicle, not, as is probably obvious, a representational content). Views that identify the former with the latter fall under the rubric of the representational strategy, giving up the idea that singular thought is fundamentally a relation to its object.

My main case against these views is the case that I make against the representational strategy generally throughout this dissertation. For the most part I will also be working at that level here. The pair of mental file approaches I consider are motivated in significant part by examples that are supposed to undermine the thesis that acquaintance is necessary for singular reference and singular thought. The resources of chapter 3, however, permit us to interpret these examples differently. Besides serving simply to defend the disjunctive strategy against the representational strategy, this is an interpretation that avoids some doubtful claims associated with the two specific accounts of singular thought in terms of mental files that we will be looking at.

4.1.2.1 Jeshion's significance condition

I'll start with a view that dispenses with acquaintance altogether. Robin Jeshion (2010) proposes that there is instead a *significance* condition on singular thought (i.e., from her point of view, on thinking through a mental file). Her intention is to stake out an alternative to an epistemic requirement like acquaintance that does not push us into the camp of semantic instrumentalists, for whom acquiring the use of a name, without any

---

[5] Scholl and Pylyshyn (1999) report empirical results that bear on this issue. Subjects tasked with tracking multiple items on a screen had impaired performance when those objects appeared to implode and explode in and out of existence. There was no impairment when the visual effect was one characteristic of occlusion ("accreting and deleting along a fixed contour" (271)). They interpret their findings as supporting the thesis that object tracking "survives interruptions of spatiotemporal continuity only when the character of the interruption supports the interpretation of an occluding surface" (287).

special requirements on *understanding* it, is sufficient for singular thought about its referent.[6]

There are two main components to Jeshion's proposal, which might be endorsed or rejected independently. One is the idea that a subject's cognitive system is what determines the nature of the thought that he entertains. This is what sets the view against semantic instrumentalism: whether one has a singular thought—which here is presumed to be a matter of which kind of representation one's cognition employs—is not under one's direct voluntary control. "When the conditions are ripe for singular thought, cognition creates singular thoughts for us and we cannot put a halt to them" (21); when conditions are not ripe, such thoughts will not be produced.

The second component of the proposal takes the form of a suggestion about how one's cognitive system determines whether to engage the representations whose employment is constitutive, on this view, of singular thought. The suggestion is that the basis of that determination is the significance of a situation for the thinker. Jeshion describes a case in which, passing a car with a parking ticket, she remarks of the unknown driver that "she's going to be upset" (11):

> Insofar as I have no interest in tracking the car owner, no affective states associated with the car owner, my mind is tailored to use the [deferred] pronoun merely as going proxy for an attributive use of a description. (21)

If, in contrast, Jeshion did have an interest in tracking the owner of the car or felt strongly about the incident, she could (as she would put it) have singular thoughts about the car owner.

Note first a worry that this proposal illustrates for the idea that we can *detach* singular thought from the object by giving up acquaintance (making singular thought an object-independent capacity) while still counting the thought as being *about* that object. If we grant that entertaining a thought with a singular content does not require that the world beyond the subject be one way rather than another, what it takes for a singular thought so understood to be putting the subject in cognitive contact with a particular thing is a further question. (On an account of singular thought as object-dependent, that question is answered by the same conditions that make the thought possible at all.) A car on the street could have several owners, or none. One can understand how a condition like significance could make one's thinking more or less object-oriented in character: focusing, for example, on patterns of property co-instantiation rather than on properties themselves, *supposing* that the car has a particular owner for the purposes of further speculations. But clearly significance does nothing to make one's mental state more determinately about one person rather than another, and this is a point that Jeshion's descriptions of her cases obscures somewhat.

---

[6] This is one expression of David Kaplan's "Instrumental Thesis," according to which new expressions not only expand our powers of reference but "broaden the horizons of thought itself" (Kaplan 1989: 603). In particular, Kaplan says of descriptive names, whose referent is determined by a description rather than a baptism, that these are "directly referential and […] have the capacity to enlarge what we can express and apprehend" (605).

A reason to be concerned about the significance condition specifically is that it excludes clear cases of singular thought.[7] Rather than trade in counterexamples, however, I want to consider *why* one might be drawn to such a condition; for I think Jeshion's examples have a different moral than the one that she herself draws.

Take a circumstance of the kind that is supposed to give rise to singular thought on Jeshion's view. Returning to a remote campsite, a camper finds her tent shredded and her food scattered—typical signs of bear activity. She becomes convinced that her campsite has been visited by a bear; driven on by isolation-induced paranoia, she also becomes convinced that the bear in question is just out of sight and waiting for its opportunity to predate her. Images of a menacing bear flash before her mind's eye, and she imagines how the attack will be described in news reports, whether the animal responsible will be identified, and so on.

Certainly psychological occurrences like these do often accompany singular thought about an individual, and I take it that this observation is what drives the suggestion that they are sufficient for the achievement of a mental state of that kind. But a much more plausible account of the resemblances between the campsite case and paradigm instances of singular thought is suggested by the discussion of chapter 3.

It must first be noted that the campsite case does not fall neatly into the two categories we considered there, "innocent" and "conniving." On the one hand, even as she swivels her head around in panicked surveillance, it is obvious to the camper that she is not acquainted with a bear (a point of contrast with "innocent" cases, in which it appears to the subject that she is having singular thoughts). On the other hand, her anxious reaction may not be under her control, whereas it's plausible that pretending or make-believe are voluntary by definition (a point of contrast with the feature of cases that inspired the label "conniving").

Recall, however, that the notion of pretence we adopted from Walton (1973; 1978; 1990) is a form of imagining that is structured around, and takes its cue from, something in the real world. We also considered that illusions of singular thought and pretended singular thought can both have their basis in informational states, psychological posits that have representational content *and* purport to derive from an individual. The wider point is that pretence and certain kinds of misapprehension are structurally similar. This feature explains the utility of pretence for the conniving user of an empty name who wishes to engage with the innocent user, the possibility of the former "getting on the same page" as the latter by pretending. The structure is also what we need to call upon to make sense of the campsite case, so it is not essential that the camper's situation qualify as straightforwardly innocent or straightforwardly conniving.

I propose to understand the case in the following way: the camper is engaged in existentially creative, prop-based imagining. Both of these features are responsible in

---

[7] (King 2020, 95). I might see a misshapen cactus on a walk, think to myself, "that's weird," and move on. Is the thought really disqualified from being singular by the fact that the cactus is insignificant to me? Supposing that the answer is "no," one way out would be to reconceive significance so as to include factors like salience or perceptual attention. Jeshion's examples—a threatening animal, an adoptee's biological parents—suggest that she really does have in mind something like "being cared about by or important to the subject," but that may not be essential to the position. A risk of the suggested emendation, however, is that it might lead us back into the territory of epistemic criteria, so in the general direction of acquaintance.

their own ways for the resemblance between the campsite case and cases of singular thought—or between the campsite case and the empty cases we have been focusing on, which the former also falls short of for reasons that I shall explain momentarily. That the imagining is existentially creative is what explains its *intrinsic* "object-oriented" character. The imaginative episode kicks off with the introduction of an object, perhaps under the guise *the cause of this mess*, to the world-as-imagined. Within the boundaries of the imaginative episode, this object accrues various properties as the episode unfolds (being a bear, having eaten the spatula, etc.). That the imagining is prop-based is what explains the extent to which it might strike us as also *extrinsically* "object-oriented," that is, having to do with a thing outside the mind. The camper's thinking is guided by the way that the world is, by actual facts about her situation.

I intend for these observations to account for Jeshion's intuition that the camper has a singular thought. Additional considerations account for what I suspect is the more widespread judgment that significance is not actually a plausible condition, necessary or sufficient, on singular thought—a reaction that is by no means the exclusive purview of those who embrace an account of singular thought as object-dependent. For while the campsite case can be brought under the umbrella of the view from chapter 3, it is also importantly different from the kinds of cases that initially appeared to threaten the object-dependent conception of singular thought. The props in those cases were informational states that *purported* to acquaint the subject with an object of thought. The props in the present case are primarily just features of the state of affairs that the camper confronts (e.g., the gashes in the tent are big, so *the claws of the bear are big*). Even if her perceptual representations play a role as well, these do not represent that a bear is present, only the aftermath of the prior presence of a bear—there is by design neither acquaintance nor the appearance of acquaintance in this scenario.

In short, an assessment of this situation inspired by Walton and Evans permits us to taxonomize the various cases at issue in a way that respects the similarities and differences between them. We have cases of singular thought, which is acquaintance-based and hence object-dependent; we have merely putative cases of singular thought, accounted for by reference to information that purports to derive from an object; and we have other cases of thinking that resemble the previous two along some dimensions, yet which are nonetheless fundamentally descriptive in character.

4.1.2.2 Recanati's acquaintance-based account

Let's now turn to François Recanati's view, which is the most thoroughly worked-out theory of mental files and their role in cognition.[8] Recanati holds that there is an epistemic condition on opening a mental file, and in the case of spatio-temporal particulars, that condition takes the form of acquaintance.[9] Problem cases like Jeshion's have been used to argue that acquaintance is too demanding a condition. (The idea being, again, that these look like cases of singular thought, ought therefore to be classified as such, and yet are not cases where the acquaintance requirement is fulfilled). Though he is basically committed to acquaintance, Recanati is moved by some of these objections to

[8] The claims I discuss here appear in (Recanati 2012).

[9] Like Evans, Recanati takes it that when you pick up a name from others this requirement can be discharged so long as someone in the chain of users was acquainted with its bearer.

attempt a compromise. The suggestion is that while opening a mental file requires acquaintance, this requirement only has the force of a *should*: the function of a mental file is to collect information from objects, and so involves acquaintance, but that function need not always be fulfilled.

This assessment flows quite naturally from the identification of singular thought with thinking through a mental file. A physical file folder can be "intrinsically" in order—it has a name on it, it contains documents that tell a coherent story—and yet "extrinsically" defective, in that the name doesn't refer. For something to function as a file in a system of information organization, it the information it collects need not actually derive from a single individual or from anyone at all. To register the defect in these cases we need to look beyond the file itself and its role in the system, to consider the use to which it is put (its function for the user or designer of the system), which is serving as a repository of information about a single, real person.

Recanati claims in particular that a mental file can be opened on *anticipation* of acquaintance. In such a case one has a "vehicle" for thinking singular thoughts about the object in question (169), though one may be as-yet unable to entertain thoughts whose contents are singular. He also speculates, in response to one of Jeshion's challenges to acquaintance conditions, that *imagined* acquaintance could be sufficient for singular thought in the "vehicle" sense (168). In addition, Recanati entertains the idea that a subject can have a singular thought ahead of acquaintance—singular at the level of vehicle and content alike—if she correctly anticipates that she will become acquainted with the referent of the term she is using.

Even if there are such things as mental files, which play the role in the cognitive economy that is usually ascribed to them, we might worry about the suggestion that it is in the function of a mental file to be acquaintance-based. There are plenty of cases where it's intuitive that a subject opens a mental file in the absence of acquaintance, either actual or anticipated, with its referent, and yet which do not strike us as deviant or defective.[10] One class of cases involves thought about fictional or mythical entities. Mental files are supposed to explain how we manage information about objects, track them across changes, and distinguish them from others. Suppose you are studying the gods of the Babylonians in school. You have information (in Evans's sense, inclusive of misinformation) coming in about various putative entities. You learn that Tiamat was the goddess of the sea, was often depicted as a dragon, and was defeated in battle by Marduk. Asked on a test whether a goddess of the sea was defeated by Marduk, you would answer "yes;" this looks like an instance of trading on identity. In short, you have standard reasons to open a mental file, and you engage in the characteristic activities that mental files are invoked to explain. One could construct similar examples involving remote sectors of the universe, or undetectable entities that we have theoretical reason to posit.

Recanati himself asserts that we use mental files in ways independent of the function he gives them—independent of actual or anticipated acquaintance. He identifies various "derived functions" that these files can serve (171). Even so, his position as stated has the consequence that such uses of mental files are *malfunctional*. This doesn't seem right as an assessment of the case just described. I propose that this odd outcome is, like

---

[10] I emphasize that in saying this I am not conceding that to cognize by means of a mental file is to have a singular thought. To the contrary, I am presenting a reason to be suspicious of recruiting mental files for the purpose of distinguishing singular from descriptive thought.

Jeshion's invocation of a significance condition, the product of incorporating some genuine observations in a suboptimal way. The case about the Babylonian gods does have a secondary status, a derivative character. Recanati is forced to account for this in terms of malfunction. On a pretence-based view, we cash it out in terms of the relationship between genuinely φing and pretend-φing: the student is engaging in some heuristically useful make-believe. In contrast to Recanati's picture, this last proposal secures derivativeness without defectiveness. It does that by giving ultimate credit to objects themselves for organizing our cognition in the object-oriented way that mental file theorists rightly emphasize.[11]

4.1.3 Conclusions

Of the roles that acquaintance played in Russell's epistemology, which are *we* committed to or in need of? Presumably not the role of ensuring that all of our referring expressions are guaranteed to refer, or the role of ensuring that the immediate objects of perception are perfectly known. Acquaintance lives on, first and foremost, as the ground of our thinking about *particular things*: as we saw in chapter 1, where I described Strawson's "massive reduplication" worry, this is a role that descriptions cannot play. It's imperative that we not confuse this positive role for acquaintance with the negatively defined role of being the contrast class for thought by description. Conceived in the latter way, acquaintance appears substitutable by other conditions—significance, for example, or merely anticipated or imagined acquaintance—that produce a mental state that intuitively deviates from paradigmatic thought-by-description in *some way or other*.

At one level there is no question that posits like mental files help to account for features of our thought and talk about things in the world around us. The potential of such posits to help us deal with empty cases in particular stems from the fact that they can fulfill certain jobs that, before we explicitly considered the implications of the possibility of empty cases for the nature of singular thought, we might have assigned to one of these things in the world. Sainsbury (2005) notes that on some direct-reference theories, "it is objects, not [representations], that are used by a subject to organize information: the mind relates 'directly' to the objects" (229). But if what information clusters around is a mental representation, a single explanatory strategy can be invoked in the empty case as in the ordinary one.

We now know, however, that the disjunctive strategy is compatible with the observation that information may be processed in the "normal" way even when there is no object of thought, no acquaintance. This is no threat to the view because the view doesn't characterize singular thought as fundamentally a phenomenon of information processing: it's a cognitive relation between a subject and something in the world, and that relation brings it about that the information coming in from the environment or from other subjects is handled in a certain fashion. In this and in the other ways I have

---

[11] Recall (§3.3) that Evans's informational states play a role in cognition analogous to the one that I've been granting, for the purposes of this discussion, is played by mental files. Following this line, one could agree with Recanati that mental files are involved in both the good case and the empty case, and give them a role in *enabling* singular thought when it does occur: what we must deny in pursuing the disjunctive strategy is that singular thought can be *identified* with thinking through a mental file (or, in Evans's terms, with being in an informational state).

described in this first part of chapter 4, the supplementation of the core claims of the disjunctive strategy by an account of the conniving use of empty singular terms as occurring within a pretence undercuts the motivations for a turn towards the representational strategy.

4.2 Intentional identity

Let's now shift our focus from attitudes considered individually to the relationships among them. Peter Geach had a talent for conjuring the kinds of sentences that keep philosophers and linguists up at night. Among these:

> (1) Hob thinks a witch has blighted Bob's mare, and Nob wonders whether she (the same witch) killed Cob's sow. (1967: 628)

Geach notices that our willingness to accept (1) does not hinge on a belief in the existence of witches. It thus illustrates what he calls *intentional identity*, wherein "a number of people, or one person on different occasions, have attitudes with a common focus, whether or not there actually is something at that focus" (627).

The phenomenon that Geach draws our attention to will figure here as both a means and an end. An understanding of the mind on which reference to things in the world has an ineliminable role in our talk about and insight into the subjective lives of others is easily abandoned when our focus is on individual attitudes; one of the lessons of intentional identity is that this picture should be reinstated. But the light that intentional identity casts on our ordinary mental state ascriptions reflects back on it in turn: an account of sentences like (1) as totally in order despite expressing no proposition just falls out of the response to Prior's challenge that I've been pursuing. I'll argue that this is the philosophical foundation that the most promising extant accounts of intentional identity judgments are missing.

4.2.1 The case of Hob and Nob

Imagine that the inhabitants of a small village are consumed by "witch mania," and that a journalist has been sent to report on the situation (628). It's in explaining how things stand with two of the villagers, Hob and Nob, that the journalist writes (1). Suppose that both the journalist and his readership know that witches don't exist.[12] The journalist's willingness to assert (1) cannot, then, be explained by a false existential belief nor by a desire to pander to that of his audience. One can nonetheless envision circumstances under which the remark would be in order, circumstances that have to do in some way with the psychologies of Hob and Nob; and this is where the problem emerges.

The special challenge posed by (1) arises from the anaphoric pronoun "she," the use of which is what expresses the (intentional) identity between Hob's and Nob's respective attitudes. Propositional attitude reports can be interpreted as transparent or opaque

---

[12] I realize that the meaning of the term "witch" is somewhat contested and that there are people who, with varying degrees of seriousness, identify themselves as witches. It's obvious that Geach's intention is for the villagers to be worked up about something that patently does not exist, and any number of mythical beings would suit this purpose; the case should be read in that light.

constructions.[13] Were the journalist simply to assert, in sequence

(2) Hob thinks a witch has blighted Bob's mare.

and

(3) Nob wonders whether a witch killed Cob's sow.

we might be satisfied to interpret each of these as *de dicto* ascriptions that do not commit the speaker to the existence of a witch. Take this line, however, and you lose the point of connection between the attitudes in (1).

Not only that, but you leave open a question that (1) settles, namely whether the attitudes of Hob and Nob are *specific* or *non-specific*. As Quine (1956) observes, "I want a sloop" has two readings: one that characterizes the speaker as wanting a particular sloop, another that attributes a desire for "mere relief from slooplessness" (177). This ambiguity is the outcome of following an intensional or hyperintensional operator with an indefinite description (certain quantifiers yield the same effect). Without further context, (2) and (3) display the ambiguity. But (1) is determinately specific: Hob and Nob are represented as converging on a single thing.

An obvious solution to both of these problems is to construe (1) transparently, along the lines of

(4) As regards some witch, Hob thinks she has blighted Bob's mare, and Nob wonders whether she killed Cob's sow.

Given that there aren't witches, however, (4) is false, so cannot be what the journalist we described intends to communicate to his audience either. In short, then, certain frontline strategies for dealing with individual attitude ascriptions seem not to be adequate to the task of interpreting intentional identity sentences.

Two further stipulations about the case will block off the most obvious avenues of escape. First, suppose that the reporter heard Hob and Nob speculate about the witch on separate occasions, and has no reason to think that the two villagers are aware of each other's attitudes. This is supposed to rule out a reading of "she" as what Geach elsewhere labelled a "pronoun of laziness" (1980, 151; originally 1962), a pronoun that goes short for a description like "the witch that Hob thinks blighted Bob's mare." Second, suppose that Hob and Nob say they first heard of the witch from the owner of the local pub, who claimed to have caught a glimpse of her creeping through the woods. This provides a plausible basis for associating the two attitudes, and one that does not hinge on Hob and Nob entertaining the same description or agreeing on the witch's properties.[14]

---

[13] Terminology from Quine (1956, 1960). Quine rejected quantification into opaque contexts; his response to the linguistic data is to distinguish two senses of attitude verbs, the *relational* and the *notional* (1956). The former relates an agent to an object of thought, the latter to a proposition (though in later sections of the paper Quine rejects propositions in favour of an extensional alternative).

[14] That two subjects each think about an object of a given kind, and ascribe to it the same set of properties, is not generally sufficient to establish that there is *one* thing both are thinking about.

4.2.2 Some possible solutions

There are evidently two sites of interest in the case Geach introduces, which may prove to be more- or less-closely related. First there is (1) itself, and other sentences like it (I leave aside for the moment the question of how that class is to be circumscribed). What is the journalist saying, and why? Then there is *what the journalist is talking about*, the Hob-Nob state of affairs; a relationship between the mental states of multiple subjects (or, as Geach notes, of a single subject at different times) that we find it natural to characterize in terms of "thinking of the same thing," even when there is no thing thought about.

Many of the accounts of intentional identity in the literature take the second of these issues as their starting point. Geach (1967) himself appeals to aspects, which are similar to Fregean senses: intentional identity is sameness of aspect. Dennett (1968) and King (1993) trace intentional identity to shared descriptions of the putative object (a move that, as already intimated, seems unlikely to fit the whole range of cases). [15] Asher (1986) suggests that making sense of belief reports calls for "a framework enabling us to speak coherently of two distinct beliefs being about the same object in an 'internal' sense so that the truth of describing two beliefs as so related does not depend on there being an *external* object that these two beliefs are about" (128). A dynamic semantics like Discourse Representation Theory (DRT) he argues, can provide the basis for just such a framework, with discourse referents representing the internal connections between sentences, and analogous devices in the mind providing internal connections between attitudes.[16] Perry (2001) has mental files related by information networks; Friend (2014) similarly argues that whether empty names "co-identify" comes down to the source of the information associated with each name. Sainsbury (2005) invokes "individual concepts," which are individuated by their origins.

Taking a different approach, authors like Salmon (1998, 2002) and Soames (2002) cite fictions as the relata of some intentional states that others would qualify as empty;[17]

Nor is it necessary. Some say the Earth is flat, some roughly spherical, and it cannot be both; but the disagreement concerns a single object.

[15] Though see (Tiskin 2018) for an account of the semantics of intentional identity sentences that treats the anaphoric pronoun as "D-type" (that is, going proxy for a description) in a way that is supposed to be compatible with Hob and Nob's ignorance of each other and disagreement about the witch's properties.

[16] Conceived near-simultaneously by Kamp (1981) and Heim (1982), the first fragment of DRT aimed to make better sense of the behaviour of indefinite noun phrases and anaphoric pronouns. It is characterized among others by its focus on *discourses*, rather than individual sentences, as the main units of analysis, and also by the space that it opens up between two aspects of interpretation: matters internal to the discourse (concerning the way that information accumulates and is organized, the links between sentences, and so on) and the matter of whether and how the content of that discourse corresponds to reality. Referring expressions are assigned discourse referents, representations that are used to build up a structure in the first phase of interpretation. The discourse thus schematized is subsequently mapped onto the world; at this point, ordinary (non-discourse) referents are assigned and claims can be assessed for truth and falsity.

[17] Soames is committed to this approach for the products of fiction, neutral on the question of products of myth (where often the intention is to make a report of fact, not create an abstract object).

Lewis (1968) makes similar use of merely possible entities. Facts about intentional identity then come down, as they do in ordinary cases of co-reference and coordination, to facts about the relationship between objects of thought: that these are identical or (as Lewis must hold) stand in the counterpart relation.[18]

Notice that these are accounts of intentional identity in the spirit of the representational and ontological strategies, respectively. We know from considering Prior's challenge that these do not exhaust our options as concerns individual singular attitudes, and so we should expect that the same is true as concerns multiple singular attitudes "with a common focus." Now, as was clear in chapter 1, the decision to adopt the disjunctive strategy does not in and of itself settle how the empty case is to be characterized. By the same token, it doesn't generate an account of intentional identity in the way that the representational and ontological strategies do. In the meantime, however, I've settled on an account of the empty case—in particular, an account *both* of what is going on with the subject whose mental states are in question, *and* of how to interpret the attitude reports we make in such contexts. We have in hand, then, the materials required to address both facets of Geach's example: what the journalist is getting across such that (1) is apt, and what it is for the psychologies of two subjects to be related in the ways that the Hob-Nob example brings to our attention.

4.2.3 An oversight

While there is not a set of views in the literature on intentional identity corresponding to the disjunctive strategy, what one *does* find is a smattering of accounts whose relationship to the ones outlined in the previous section is not immediately clear: views whose unifying thesis is that the literal meaning of (1) is not the content at issue in the journalist's report.[19] The disjoint character of the familiar range of options stems from the fact that, as we've seen, the accounts in terms of representations and "exotic" objects are primarily addressed to the psychological question that Geach's example raises; whereas this last cluster of views is primarily addressed to the linguistic one.

The representational and ontological treatments of the Hob-Nob case identify the *fact* in virtue of which Hob and Nob stand in the supposed intentional identity relation. Presumably the journalist is talking, in some more or less direct way, about that fact; whether he does so thanks to the literal meaning of (1) or thanks to what (1) conveys, implies, etc. in the context is not especially to the point. That is: whatever these views identify as the relevant relationship among mental states, translated into a content that an ordinary person might plausibly entertain (the journalist presumably does not *mean* by (1) that Hob's and Nob's mental files are appropriately related, or that they both think of possible witches that stand in the counterpart relation), it could be offered either of an analysis of (1) or as a spelling out of the message that it conveys pragmatically. So it would be odd if the mere fact of insisting that (1) is not to be taken literally were sufficient to generate a distinct kind of response to the problem of intentional identity in the way that, if we glance at the existing dialectical landscape, it appears to do.

---

[18] On the counterpart relation, see (Lewis 1968), (Lewis 1986).
[19] (Crimmins 1998), (Manning 2015), (Hoek 2018).

Looking at some of the claims made by defenders of the third cluster of views will make it clear why the literature has the shape that it does. Mark Crimmins (1998) and Luke Manning (2015) both tie intentional identity to the problem of negative singular existentials that I alluded to in §3.4—Manning's central claim is that the problem of intentional identity can be *reduced* to the problem of negative existentials. Recall that this latter problem is what led Evans to draw on Walton's make-believe theory, and that both authors propose solutions to it along broadly similar lines. (I've suggested that the primary interest of the appeal to pretence lies elsewhere, in what it allows us to say about empty attitude ascriptions.)

Crimmins turns to Walton and Evans in this connection; in particular, he pursues their suggestion that we might identify a real-world truth condition for some sentences containing singular terms that lack a referent by means of a principle linking truth-in-the-pretence with truth simpliciter (1998, 4). Addressing the Hob-Nob sentence (1), Crimmins suggests that the "real-world requirement for fictional truth, and also the truth condition of the serious statement" is the following:

> There is a witch-generating mode of presentation *m* such that Hob believes: [*m*] has blighted Bob's mare, and Nob wonders whether: [*m*] killed Cob's sow. (40)

He characterizes this as a "fictional-object" reading of (1).

Crimmins's proposal has some features in common with the kinds of views that I have associated with the representational strategy and the ontological strategy. It resembles the first in the way that he is conceiving of modes of presentation, notably as object-independent. Evans would deny that there is a mode of presentation in the empty case, just as he denies that empty singular terms are meaningful. Crimmins is working with a notion of a mode of presentation, evidently not the same one as Evans employs, on which these are available to be leaned upon in the explanation of the empty case. Next, Crimmins's proposal resembles the ontological strategy in its invocation of fictional objects. There is of course a sense in which fictional objects can be said to figure in the pretence-based account of empty cases I introduced in chapter 3: there are objects in the fiction, in the pretence. But one must be careful not to elide the substantial difference between this claim and the one characteristic of some developments of the ontological strategy, which means by "fictional object" a kind of real entity, an abstract artefact generated by authors, pretenders, and (by some accounts) also by the thoughts and reflections of people who are just confused.

That said, we can make a few adjustments to bring this suggestion into conformity with the way of elaborating Evans's views that I presented in the previous chapter. A role like that assigned by Crimmins to modes of presentation belongs to informational states. These are the representations that are common to cases of singular thought and cases of merely-putative singular thought. The objective fact that we can point to in connection with the subjective reality that we use (1) to express is that Hob's and Nob's information has a common source (they are both taking up the testimony of the bartender, whose sensory informational state may have kicked off the whole affair). The reason neither of them is in a position to have singular thoughts, and the reason why this is not just an ordinary case of two people thinking about the same entity, is that if we trace their information to its origin we find that it does not originally emanate from a suitable object

(from a witch).[20] Hob and Nob don't realize this; their situation is indistinguishable to them from a situation in which they have information-based, singular thoughts about an entity that the bartender genuinely saw. That is the further, subjective fact that we need to engage in pretence in order to capture.

On Crimmins's presentation, the main payoff of the use that Walton and Evans make of pretence is that it provides us with a roundabout way of securing a truth condition for sentences that, like (1), seem not to be analyzable in a way that captures their intuitive significance. This is another point at which I would part ways with his account of things. The idea that sentences like (1) are true *simpliciter* if they are true in the pretence, though one that Walton and Evans both endorse,[21] obscures the thoroughly novel character of the strategy that their ideas enable us to pursue.

On a standard pragmatic account, the journalist is justified in uttering a sentence that is strictly speaking meaningless (or as some would instead hold, false) because he thereby succeeds in getting his audience onto a *true proposition* about the mental states of Hob and Nob. This cannot be a proposition about their attitudes towards a concrete entity (the sort of objects that they presume witches to be), because there is no such object. Notice that on this line of thought we are forced into an internalist conception of singular attitudes just by our efforts to explain why a certain kind of utterance is appropriate! Situating the journalist's utterance in a pretence, we can preserve the obvious analysis of (1) as

> As regards some witch, Hob thinks she has blighted Bob's mare, and Nob wonders whether she killed Cob's sow.

and maintain that singular thought is object- and belief-dependent, as Evans suggests. We can resist what is strongly implied by Geach's definition of intentional identity (when several attitudes "have a common focus, *whether or not* there is anything at that focus"), namely that, in the terms of his metaphor, *having a common focus* is prior to *having a common object*. The reality is, as that metaphor itself suggests, the other way around: if we can make sense of attitudes having a common focus, it's because there is something they are focused *on*. The illusion of a unified, object-independent phenomenon of coordination arises when we fail to recognize that superficially similar cases in fact belong to fundamentally different kinds, one of which we understand by reference to the other. There are cases where multiple attitudes are about, on non-descriptive grounds, a single entity; and then there are cases, like that of Hob and Nob, that we handle by pretending that they are cases of the former object-involving type.

What is missing in Crimmins's view and in the literature on intentional identity more broadly is a connection between the claim that (1) is not literally true and the

---

[20] One of several ways that a subject *S*'s information-based thought about an object *O* can be, Evans suggests, "ill-grounded" (1982, 132). (i) *S*'s information concerns *O*, but her purported mode of identification singles out no object. (ii) *S*'s information concerns *O*, but her purported mode of identification singles out some object *O** distinct from *O.* (iii) *S*'s information concerns no object, but her purported mode of identification singles out *O**. (iv) *S*'s information concerns no object, and her purported mode of identification singles out no object. Visual learners may wish to consult the accompanying diagrams (133).

[21] See (Walton 1990, 399) and (Evans 1982, chap. 10, sec. 10.3).

disjunctive strategy. The representational and ontological strategies are naturally allied with an account of the journalist's report on which it expresses or conveys a true proposition. The disjunctive strategy is naturally allied with an account of that report on which it is apt precisely because there is no proposition (external to the pretence) that it expresses, since the kind of proposition required to do justice to the relationship between Hob's and Nob's attitudes, as well as to their individual perspectives, is not available.

These alliances are not *compulsory*: the suggestion that there is no proposition expressed by (1) belongs to the particular spelling-out of the disjunctive strategy that I am defending, not to the essence of that style of response to Prior's challenge. Overall, there are six main ways of dealing with Geach's example, according to how one answers the following two questions:

- Does (1) communicate a true proposition? (yes/no)
- What is the right account of Hob's and Nob's individual attitudes? (representational strategy/ontological strategy/disjunctive strategy)

The fact that several of these combinations are not represented in the literature explains the impression of disjointedness that I alluded to above.

4.2.4 The lesson of intentional identity

We might supplement Geach's story about Hob and Nob with an invented tale of our own. Imagine a second village, whose residents are distinguished by their pragmatic and straightforward approach to life. They are grounded in the concrete and common-sensical, not especially prone to abstract reflections.

Suppose that these villagers do not have a concept of mental representations as such.[22] They acknowledge subjects of experience, understood just as occupiers of perspectives on the world; concrete particulars; and relations that subjects can stand in to these particulars. Perceiving is being aware of the environment. Having a pain is being aware of the condition of your body. When something has been perceived, it (itself) becomes a potential object of various cognitive relations: wanting, fearing, contemplating, etc. For these villagers, there is no seeing without something seen, no wanting without something wanted, just as there can be no eating without something eaten.

Let's focus in on one of these people, whom I'll call "Fub." Like Geach's journalist, Fub also makes a visit to the village of Hob and Nob (though not because of the witch mania; Fub is opposed on principle to this sort of spectatorial behaviour). For the most part, Fub has no trouble. He encounters cases of veridical perception, of thought and talk about things that exist. He encounters groups of people who are attending to or talking about a single, spatio-temporal entity. He encounters people acting on things that are there to be acted upon.

Fub constructs for himself a little model of this neighbouring village, listing out the people and objects to be found there and the ways that they are interrelated. This model serves a dual purpose for him: the whole provides an account of how the village is,

[22] My descriptions here are loosely inspired by Butterfill and Apperly's (2013) articulation of what they call a "minimal theory of mind."

but parts of *that* account also constitute accounts of the subjective life of a particular villager. For instance, there is Hob, and the village clocktower, and Hob standing in the "sees" relation to the clocktower at such-and-such time. That Hob sees the clocktower is not just a fact about the state of affairs in the village; it also records a subjective fact, a sliver of what it is like to be Hob in that moment, because the clocktower is a constituent of his sensory experience. Or there is Nob, and *that horse* in the square, and Nob standing in the "wants" relation to that horse. That Nob wants the horse is not just a state of affairs in the village; it also records a subjective fact, because by identifying an attitude and what it is directed toward Fub has given what is for him a full characterization of (what we would call) an intentional mental state.

Often Fub can tell, just by looking at a scene or assessing a villager's behaviour, what they are sensing through which perceptual modality, or what attitude they have towards which object. But he prizes above all else the villagers' own sincere first-personal reports, because they themselves occupy the perspectives that he is interested to describe and understand. One day, Fub again comes across Nob, who (Fub notes) is squatting down and looking at the ground, even though there is nothing there. Fub asks Nob what he is doing. Nob replies, completely seriously: "I am admiring this magnificent butterfly."

Now Fub faces a challenge, and a choice. All he has to work with in his understanding of the mind are things and relations to them. Here there is no thing and no relation—no butterfly, no admiring of it. At the same time, Nob is the authority on his own conscious life, that is, in matters of describing *the world from his point of view*. Must Fub simply give up on Nob—on making sense of him, on talking to him, on reasoning with him? If the people in Fub's village, like most (non-fictional) humans above the age of two, have a capacity to engage in make-believe play, the answer is "no." Fub has the resources he needs to engage with Nob, if he so chooses. Without becoming any less aware of what is in fact the case, he can *pretend* that Nob is giving voice to his perspective on the world (where, again, Fub understands "having a perspective on the world" in the most flat-footed possible interpretation of that phrase). This does not require that Fub think of Nob as representing the world as being a certain way. It does not require him to expand his existing ontology by coming up with the idea of an abstract object, which can take over the role in his subject-relation-object schema that, were the case a normal one, a butterfly would have played.

I use this fanciful description to bring into relief features of our own way of operating that are easy to miss; that can get drowned out, in particular, by philosophical preconceptions about the mind and about the workings of language alike.[23] The case that Geach brings to our attention is not a puzzle to be resolved within the frame of the representational strategy, but a counterexample to that now-prevailing perspective on the relationship between the mind and the rest of the world. Restricting our focus to individual attitude reports, it's easy to conclude that their relational syntax has no deep significance. Intentional identity sentences reveal the appearance that objects have an essential role to play in our characterization of certain mental states—not 'intentional objects,' construed in some mysterious fashion, but *things*—to be surprisingly robust.

---

[23] I say more about the latter—making explicit a point that I have so far only gestured at—in chapter 6 (§6.1).

From the perspective of the disjunctive response to Prior's challenge, however, this should not be a surprise at all.

Chapter 5: From Thought to Perception

In chapter 2, I highlighted a feature of our use of such verbs as "thinking of," "wanting," "loving," and so on: whereas we have a concept for an epistemic state that falls short of knowing (belief) and a concept for an experiential state that falls short of perception (hallucination), there is not an overt, linguistically-marked contrast between ascriptions of singular thought and ascriptions of states that fall short of it. We've now had the opportunity to observe how this feature contributes to the impression that a single kind of mental state is at issue whether the object of thought exists or not; we've also seen it give rise to confusion, in the section on intentional identity in the previous chapter.

This phenomenon is not unique to the case of thought. For while it's true that we can describe a subject as *perceiving something* or as *hallucinating something*, we also use the verbs associated with each of the sense modalities—"sees," "hears," "feels," "smells," "tastes"—to characterize veridical perceptions and hallucinations alike. In this chapter, I extend the account of our ordinary mental state ascriptions that I introduced in chapter 3 to demystify the possibility, which has generated almost a century of debate, of using the verbs of perception in these two ways. I begin by introducing the problem and previewing my solution to it (§5.1). I then provide an overview of the existing literature, which I divide into two main phases: the early debate (§5.2) and its contemporary revival (§5.3). §5.4 is diagnostic in character, explaining where I take the debate to have gone wrong. In §5.5–5.9, I gradually make my case against the usual explanations of the use of perceptual reports at issue. §5.10–5.13 set out my account of this use and indicate how the debate should be resolved.

5.1 Orientation

The initiating observation of the debate about the verbs of perception is not simply that there are two uses of perceptual verbs, but that, on further reflection, those uses seem like they might be incompatible. Here's why. We can first note that reports typically imply the existence of the object said to be seen, felt, tasted, heard, or smelled. If you are on the phone with a friend who tells you that they smell smoke, for instance, the default assumption is that there is smoke where they are. If someone searching for a dog exclaims, "I see her!" we presume that he has found the dog and has not merely been struck by a vision. Most decisive are exchanges like the following:

A: I heard a voice! B: You can't have heard a voice, there's nobody in there.
A: Will we see any whales? B: No, they're further south at this time of year.

In such cases, the presence of an object is treated as a condition on the possibility of perceiving it.

But now suppose that, having ingested a dose of psilocybin, you see a persimmon on the table in front of you—or so you are inclined to say. Your sober companions see none, and indeed there are no persimmons around. It's natural to introduce cases of hallucination using a perceptual verb, and you might subsequently relay the incident in

like fashion: "I saw a persimmon of the most vivid orange…" Yet in the envisaged case there was no persimmon on the table to be seen, and you know this.[1]

There are two main hypotheses by which to account for these facts in the literature on the verbs of perception. The first is that perceptual reports are ambiguous: one literal interpretation is at issue in reports of veridical perception, and another in reports of hallucination. The second is that we sometimes use these reports—which strictly *do* imply the existence of an object perceived—to convey by pragmatic means a description of the subjective character of an experience.[2] The longevity of the contest between these two positions is, I claim, no accident: there can be no resolution via either of the familiar routes, because both are premised on the same false assumption. Parties on either side of the debate take it that perceptual reports are appropriately invoked to talk about hallucinations because their use gets across, by some means or other, a proposition with the following two features:

> *(i)* its truth does not require the involvement of a thing perceived, and
> *(ii)* it provides an accurate characterization of the subject's perspective.

I'll argue that no proposition can meet both criteria. That being so, the interesting question raised by the way we talk about anomalous sensory experiences is *not* whether such talk is literal (the answer to that question will be a straightforward "no"). What's puzzling is how and why we persist in talking about the things we see when we know that we aren't seeing anything. As I interpret this behaviour, it buttresses the claims I have already made in the context of singular thought: our grasp on the experience of the perceiving subject is utterly dependent on our grasp on *what is perceived*, where the "what" is just ordinary things, events, and states of affairs in the world. Since these materials are absent in cases of hallucination, we are prompted to engage in existentially creative make-believe whose development is dictated by the sensory representations of the subject whose experiences we are interested in. The perceptual reports that have provoked so much discussion occur within the confines of this activity and must be interpreted in that light.

---

[1] There are accounts of hallucination as involving awareness of something. Sense-datum theory is the classic example, but see also Byrne and Manzotti's (2022) suggestion that to hallucinate is to be aware of a physical object—whether an ordinary thing perceived in the past, or a "gerrymandered object" that combines remembered ordinary things or their parts in various ways (341). On that view, there may after all be a persimmon that you are aware of when you hallucinate. I ignore this possibility because there's no reason to think that we *mean* to credit hallucinating subjects with present awareness of past objects of perception or gerrymandered objects, even if that is the right account of the metaphysics of hallucination.

[2] Some version of the first approach is defended or assumed by Moore (1905–6, 1952) Ayer (1940, 1956), Smythies (1956), Hintikka (1969), Lewis (1983), Harman (1990), Chomsky (1995), Brogaard (2015, 2018), and Bourget (2017, 2019). Anscombe (1965) is often placed on this side of the debate, though I doubt that her position can be so neatly classified. Advocates of the second, pragmatic approach include Austin (1962), Dretske (1969), Grice (1989), Soames (2003), and Hoek (2018). D'Ambrosio (2022) reports empirical evidence for the intensionality of "to sense" and "to perceive," but not "to smell," "to see," "to hear," and the like.

5.2 The early debate

The first surge of interest in the meaning of the verbs of perception is centered at Cambridge and Oxford at mid-century; its main players are Moore, Ayer, and Austin. The debate originates in controversies about the existence and nature of sense-data. Among the points of dispute: whether such entities already figure in our pre-theoretical understanding of experience. It's in this connection that Moore directs us to a scene from Shakespeare:

> When Macbeth says 'Is this a dagger which I *see* before me?' the sense in which he is using 'see' is that which I am now calling 'directly see,' [...]. Macbeth is represented by Shakespeare as having directly seen an 'object' (though, of course, not a *physical* object, any more than an after-image, seen with closed eyes, is a *physical* object) to which he referred by the word 'this,' and about which he asked 'Is *this*, which I see before me, a dagger?', and we all understand perfectly well what *sort* of experience Shakespeare is representing Macbeth as having had. (1952, 630; emphasis original)

Moore is sometimes characterized as holding that there is a sense of "see" that does not imply the existence of what is said to be seen; a sense in which Macbeth can truly say that he sees a dagger, though doubting that there really is one in front of him. But this misrepresents both Moore's position and the quotation that he presents as evidence. What strikes Moore is that Macbeth seems to make a demonstrative reference to something ("this") and then ask whether *it* is a dagger, or merely a product of the heat-oppressed brain. Thus interpreted, the line is congenial to Moore, because the obvious candidate for the referent of the demonstrative is a sense-datum: a private, non-physical *something* with a certain shape and colour, which can be present to the mind regardless of whether one veridically perceives or hallucinates.[3]

What inspires Moore's claim about "see" is thus an observation, not initially about the verb itself, but about the range of objects that we claim to see. We habitually speak of seeing daggers, tables, and other physical things; we occasionally speak of seeing afterimages and other arrays of sensible properties that lack an environmental basis. If the latter statements are interpreted as reports of encounters with sense-data, and sense-data are a common element of veridical perception and hallucinatory experience, then physical objects and sense-data are not seen in the very same sense, or in the same way. Rather, we (indirectly) see physical objects by (directly) seeing sense-data; and Moore finds evidence in Shakespeare for a corresponding ambiguity in the perceptual verbs as we already use them.[4]

---

[3] I'll often return to Moore's example for the sake of continuity, but note that, since we are interested in whether the sentence is literally true, it introduces a potential confound: Macbeth, though based on a real person, is just a character in a play. The reader should ignore this complication in what follows, and treat the name as straightforwardly referring to a non-fictional subject of experience.

[4] One might construe the suggestion here in the manner of Frank Jackson (1977), who, in discussion of Moore's view among others, seems to treat "see" as univocal and attribute its contrasting uses to modifiers that are usually unarticulated in ordinary speech: e.g. "mediately," "immediately," "directly," "indirectly" (6–7). Nothing much will hang on this interpretive

Ayer opens *The Foundations of Empirical Knowledge* (1940) with a reconstruction of the argument from illusion, which moves from the observation "that material things may present different appearances to different observers, or to the same observer in different conditions" to the conclusion that we only ever perceive sense-data (3). To Austin's consternation, Ayer presents the sense-datum theorists as simply proposing to use perceptual language in a new way:[5]

> ...what the advocates of the sense-datum theory have done is to decide both to apply the word 'see' or any other words that designate modes of perception to delusive as well as to veridical experiences, and at the same time to use these words in such a way that what is seen or otherwise sensibly experienced must really exist and must really have the properties that it appears to have. (24)

This amounts to a denial that Moore's "directly see" is already in circulation. Ayer nonetheless affirms that the perceptual verbs have two contrasting applications, and offers his own account of these.

Using "see" in its most commonplace meaning and without negation, Ayer claims, the truth of what one says requires "that what is seen should really exist, but not […] that it should have the qualities that it appears to have" (23). One would be using "see" in this way if, looking at the Müller-Lyer illusion, one said "I see two lines of equal length." What it takes for this statement to be true is that there be two lines of equal length and that one see them—and not anything about how those lines *appear*. Thus one speaks truly, using "see" in this first sense, despite the fact that one line looks longer than the other.

On Ayer's proposed second sense, in contrast, "it is not possible that anything should seem to have qualities that it does not really have, but *also not necessary that what is seen should really exist*" (23; my emphasis). This sense, Ayer suggests, has a close association with talk explicitly concerning appearances, and can be paraphrased away with the aid of locutions like "appears to be" or "seems" (*ibid.*). It's with this sense of "see" that we can truly say, again confronting the Müller-Lyer illusion, "I see two lines of *different* lengths." The two lines before one do not meet this description, and so, Ayer thinks, what one claims to see does not exist. This doesn't impugn the truth of the statement, because making a claim about what the subject is perceptually related to is not the point: the speaker is using the perceptual verb to describe how things look, never mind how they are.[6]

---

question; though note that Moore does for instance speak of the "sense [of 'see'] which I have called 'directly see'" (1957), a remark that suggests he has in mind two senses for the verb which he only uses the modifiers "directly" and "indirectly" to disambiguate.

[5] This claim shows the influence of Carnap's view, encapsulated in his "Principle of Tolerance" (1937, §17), that what look like substantive metaphysical disagreements in mathematics are merely reflections of different choices about how to use language.

[6] Notice that Moore's second sense of "see" combines Ayer's two conditions: one speaks in such a way that what is seen cannot but exist and cannot fail to be other than it appears. (As befits our epistemic position vis-à-vis sense-data). As we'll find, these are the same conditions that someone like Austin claims govern the use of "see"—unsurprising when one considers the sense-datum theorist and the direct realist's shared commitment to the claim that perception is fundamentally relational. How, then, can Moore and Austin be disagreeing about the language of perception? They

In *Sense and Sensibilia*, Austin takes up the mantle of the opposition to the "two senses" proposal. His main strategy is to discredit the various examples that Ayer offers as evidence for a second sense of "see," a sense without existential commitment. For instance, Ayer points out that one can say, "I see a stick that looks crooked" (e.g., observing a straight stick in water) without implying that anything really instantiates the property in question. Austin concurs that the whole sentence lacks this implication, but proposes that this be chalked up to the meaning of "looks crooked": after all, "whatever views we may or may not have about senses of 'see', we all know that what looks crooked may not really *be* crooked." (1962, 88) To the possibility of reporting that one feels pressure in one's leg though the leg has been amputated, Austin rejoins that we have as much reason to posit that "pressure in my leg" has two senses as to propose two senses of "feel"—which is to say, from his perspective, little reason at all. And whereas Ayer describes the use of the perceptual verbs to characterize (mere) appearances as a "perfectly correct and familiar usage" (21), Austin qualifies it as exceptional. We speak this way "faute de mieux" (91), managing to make ourselves understood only because our interlocutors are "aware of the *special circumstances*" (*ibid.,* emphasis original).

This early debate faded away only to be reprised decades later; we find some first hints as to why the problem was not resolved the first time around in a second debate about so-called "perceptual identification" that initially unfolded in parallel with the one we have just surveyed. [7] This disagreement spawned from the same remarks of Moore's already quoted. Returning to the passage that opened this section, we can notice that it presents Macbeth's question twice over, with a change in emphasis:

> "Is this a dagger which I *see* before me?"
> "Is *this*, which I see before me, a dagger?"

The first serves to introduce us to the (alleged) second sense of "see." The latter underlines a feature of the case as Moore conceives it: the presence of a *something* that Macbeth can get a grip on and refer to without yet knowing whether he veridically perceives or hallucinates, something that may or may not prove to *be* a dagger.[8]

Posed in the second fashion, Macbeth's question invites us to consider the different ways that we describe or identify what we see. This theme recurs in Ayer's discussion of "see," in connection with the following case:

---

disagree about the kinds of contexts where a sense of the verb "see" governed by these two conditions can be used to say something strictly and literally true. For Moore, "I see a red circle" can be true when I am hallucinating, *even though* it commits me both to my experience having a certain character and to the existence of something red and round, because to hallucinate is to be aware of sense-data that instantiate sensible properties.

[7] In addition to (Ayer 1940) and (Austin 1962), the main contributions are (White 1958, 175–183), (Hampshire and Strawson 1961), and (Hinton 1967a).

[8] There is a confusion here that only grows when Hampshire and Strawson pick this topic up from Moore. Depending on the circumstances, Macbeth may be talking about a dagger or talking about a hallucination. But there is no entity he is aware of that could either *be* a dagger or *be* a mere product of the heat-oppressed brain. Given Moore's commitments, what he ought to say is that Macbeth is aware of a sense-datum that could have come about in the usual manner (looking at a dagger) or in an unusual manner (overheated brain). In either case, Macbeth directly sees a sense-datum; in the former, he thereby indirectly sees a dagger.

> …a man will say that he sees a distant star which has an extension greater than that of the earth, but if he is asked to describe what it is that he is actually seeing, he may say that it is a silvery speck no bigger than a sixpence. Now, since it is impossible that the star should both be bigger than the earth and at the same time no bigger than a sixpence, one is tempted to conclude that one at least of these assertions is false. (22)

The descriptions that occur in these reports are, on the face of things, incompatible; and yet it seems like both can be true in one and the same situation. For Ayer, this is further evidence that there are two senses of "see." A star viewed from earth *looks* tiny but *is* massive. Selecting between the two senses, one selects between two commitments: a commitment to a certain characterization of the sensible appearances (as with "I see a tiny speck"), or merely to the existence of what is perceived (as with "I see something larger than the Earth").

Austin takes a surprising line at this stage.[9] One would expect him to dismiss "I see a silvery speck" as yet another extended use of perceptual language. Instead, however, he treats the utterance as straightforwardly true on his own favoured and univocal interpretation of the verb, namely as entailing the existence of the object said to be seen:

> If a man says 'I see a silvery speck', *of course* he 'implies' that the speck exists, that there is a speck; and if there is *no* speck in the region of the night sky at which he is looking, if that part of the sky is perfectly blank, then of course he does *not* see a silvery speck there. (94)

Austin denies that there is even an intuitive conflict between "I see a silvery speck" and "I see a massive star," insisting that "the silvery speck *is* a star" (92). He thus classifies the example as one instance of a familiar and, by his lights, anodyne phenomenon:

> …what we "perceive" can be described, identified, classified, characterized, named in many different ways. If I am asked 'What did you kick?', I might answer 'I kicked a piece of painted wood', or I might say 'I kicked Jones's front door; both of these answers might well be correct; but should we say for that reason that 'kick' is used in them in different senses? Obviously not. What I kicked--in just one 'sense', the ordinary one--could be described as a piece of painted wood, *or* identified as Jones's front door; the piece of wood in question *was* Jones's front door. (98)

Contrary to what Austin claims, however, these two cases are *not* analogous. "Jones's front door" and "a piece of painted wood" may both be, on a literal interpretation, unproblematically true descriptions of a single thing at a single moment in time. The same does not hold of being no larger than a sixpence and much larger than the earth. A single object simply cannot meet both of these descriptions at once.

That something has gone awry is suggested by a footnote of Austin's, albeit one whose content he himself takes in stride:

---

[9] As A.N. Prior also emphasizes; see the posthumous reconstruction of his draft review of *Sense and Sensibilia* by van Hulst and Cresswell (2016).

> It doesn't follow, of course, that we could properly say, 'That very large star is a speck.' I might say, 'That white dot on the horizon is my house', but this would not license the conclusion that I live in a white dot. (99n1)

We can say "that speck is a star," Austin admits, but not "that star is a speck." Well, why should this be? There is no problem going back and forth between "Jones's front door" and "a painted piece of wood." Those who pursued Austin's thought further found themselves wondering whether an as-yet-untheorized sense of "is" was at work in such sentences as "that silvery speck is a star," which would explain the possibility of going in one direction and not the other.[10] From two senses of "see" to two senses of "is": the kind of outcome Austin wished to avoid has only been kicked down the road.

Besides the bizarre consequences and the unexplained break with his previous strategy, the deep issue with Austin's response is that it causes him to overlook the lesson of this last example—of the possibility of saying "I see a silvery speck" to report what one knows to be a *veridical* perception of an enormous celestial body. It's not surprising that a philosopher with Austin's commitments would be set on edge by the suggestion that episodes of seeing can occur even in the absence of a confrontation with an appropriately propertied environmental object. But the case of the star shows us that there is no *essential* connection between the use of perceptual reports that we are trying to make sense of and such aberrant occurrences as hallucination and illusion (contrary, also, to what the subsequent literature on the perceptual verbs suggests). Anscombe (1965) emphasizes on behalf of the sense-datum theorists that "[t]here is such a thing as simply describing impressions, simply describing the sensible appearances that present themselves to one situated thus and thus" (173). The practice is in fact so widespread and so ordinary that Austin, having let his guard down in response to the fact that there *is* an object seen in the star case, doesn't recognize "I see a silvery speck" as an instance of just the sort of talk that he has otherwise made a point of militating against.

For the same reasons, Austin would not have been better off if he had classified this, too, as an extended use of language; he would simply have erred in a different way. In his eagerness to explain away Ayer's examples, Austin concedes too much to his opponent. He accepts that these uses are a special problem for his own preferred account of the meaning of the perceptual verbs: that those verbs pick out relations between a subject and some constituent of our shared environment (i.e., not a sense-datum). Austin must indeed deny that the sentences we've been considering ("Macbeth sees a dagger," "I see a silvery speck," etc.) are, on the use in question, literally true. As I have emphasized in previous chapters, however, a sentence can be false or even meaningless in a context and yet be perfectly apt—speaking in that way may even be required for the fulfillment of one's communicative aims.[11]

In making an explicit invocation of this possibility, we are faced with a question that Austin for the most part sidestepped. If perceptual reports are not strictly and literally true as employed to characterize hallucinations, illusions, and what we might

---

[10] (Hampshire 1964), (White 1958). (Hinton 1967a, 422–423) summarizes and discusses the relevant claims.

[11] This is a point that I've made in connection with the claim that some of our subjective states are not amenable to literal description. It can also be supported independently (see §6.1).

hyperbolically term "misleading" appearances (e.g., the way stars look from Earth), why do we regularly turn to them for this purpose? Answering that question requires us to engage carefully with the use that Ayer highlighted and, I would suggest, quite faithfully reported. Doing so will eventually lead us to a surprising discovery: it's through a proper appreciation of the very use Austin tried to downplay that one finds the best support for his position on the literal meaning of the perceptual verbs.

5.3 The contemporary debate

The debate about the verbs of perception has always been tied up with the metaphysical commitments of its participants. At the same time, the philosophy of perception has undergone significant shifts since the 1950s. Sense-datum theory has ceded its place, as the main foil to naïve realism, to intentionalism: a view on which perceptual states, like propositional attitudes, are fundamentally relations to representational contents. The debate about the perceptual verbs has revived subsequently to this transition. And much as Moore hoped to find intuitive support for sense-datum theory in a passage from *Macbeth*, contemporary efforts aim at buttressing the representationalist's assimilation of the perceptual to the psychological by means of claims about everyday linguistic behaviour.

The core issues and the basic positions in this second phase of debate are not new. An optimistic read on this situation is that the replacement of sense-datum theory by intentionalism is what we need to finally answer the question that Moore introduced to the philosophical consciousness. More pessimistically, however, one might worry about the risk of falling into the same traps again without even noticing—those traps having been disguised by a new theoretical dressing. My assessment aligns more closely with the latter characterization. In the rest of this section, I describe details of the contemporary debate that I will ultimately treat as peripheral, simply to provide the reader with a sense of the current literature. I proceed to an evaluation of the debate as a whole in the section (§5.4) that follows.

5.3.1 Semantic accounts

Anscombe (1965) charges the sense-datum theorist and the ordinary language philosopher alike of "failing to recognize the intentionality of sensation" (171). What everyone does recognize, of course, is the intentionality of the *mental*. The implication of Anscombe's remark is that the peculiarities of perceptual and psychological language have a common source, and so are subject to a common explanation. That idea comes to the forefront in the contemporary debate. This phase is spearheaded by Brogaard (2015, 2018) and Bourget (2017, 2019), both of whom maintain that the perceptual verbs at least sometimes produce intensional contexts, just as the attitude verbs are widely held to do.[12]

Recall from §2.1.1 that intensional contexts are standardly distinguished by failures of existential generalization and substitution and by the specific/non-specific ambiguity. We've noted already—this is the observation that kicks off the debate—that there are

[12] Since nothing will hang on the distinction, I group hyperintensional and intensional contexts together, for the purposes of this chapter, under the heading "intensional."

uses of perceptual reports which do not imply the existence of an object. Moreover, as Anscombe observes, perceptual reports are often intended to supply a "*description under which*" the subject of the ascription conceives of the object of her state (166; emphasis original).[13] Success at this task seems to be what drives our assessment of the aptness of certain perceptual reports, rather than the existence or actual character of their object; so we have evidence of substitution failure.[14] One might further expect that, when the verbs of perception are invoked to talk about hallucinations, there can be a non-specific reading of the object phrase—after all, there need be no particular in the world that causes (in the normal way that perceptions are caused) or corresponds to what is represented by such experiences.[15] In short, at least some perceptual reports show signs of intensionality.

The proposal that the verbs of perception create intensional contexts faces an initial obstacle at the level of implementation. To see why, first consider:

(1) Lilia believes that someone is outside.
(2) Lilia believes that the mayor is outside.[16]

Both of these sentences might be true even though nobody is outside; Lilia might believe that the mayor is outside without believing that Stevie is outside, even though Stevie is the mayor; and the "someone" of (1) can be read as specific or non-specific. (1) and (2) are propositional attitude reports, and in such reports the presence of all three marks of intensionality can be straightforwardly explained by a scope ambiguity. (1) and (2) each have two readings, as follows:

(1) *a.* $\exists x$(Lilia believes that x is outside).
b. Lilia believes that $\exists x$(x is outside).
(2) *a.* $\exists x$(x is the mayor and Lilia believes that x is outside).
b. Lilia believes that $\exists x$(x is the mayor and x is outside).

On the readings where the existential quantifier takes wide scope—(1) *a* and (2) *a*—the sentences entail that there is someone about whom Lilia has a belief. (1) *b* and (2) *b* do not. We can see from (2) *a* that the wide-scope reading also permits the substitution of co-referring expressions. What makes it true is that x has the property of being the mayor, not Lilia's believing that x has this property. Finally, (1) *a* and (1) *b* capture the specific and non-specific readings, respectively, produced by the quantifier ("someone") in (1).

The reader may notice that the perceptual reports with which I opened are not propositional in form. Instead, they have their closest analogue in so-called "objectual" attitude reports, reports like, "Lilia dislikes the mayor." Objectual attitude reports raise a problem for the treatment of intensional contexts just outlined. The reason is that while they give rise to the same three features, they cannot obviously be subjected to the same

[13] Recall discussion of Anscombe's claims about attitude verbs in §1.3.3 above.
[14] See also Brogaard (2018), chapter 6.
[15] This expectation turns out not to be fulfilled, as I explain in §5.5.
[16] Note that my setup of this point will involve shifting between two uses of definite descriptions that are often contrasted, one predicative and the other (arguably) referential (see e.g. Strawson 1950). The questions this raises about the semantics can safely be set aside for present purposes.

analysis as their propositional counterparts.[17] A semantic account of the workings of perceptual reports needs first to address this issue, and there are two prominent ways of doing that (Brogaard pursues one, and Bourget considers the prospects for both). Each give rise to their own challenges.

According to *propositionalists*, any apparently objectual constructions that exhibit one or more of the three marks of intensionality are disguised sentential-complement constructions. [18] the contrast I've drawn between objectual and propositional attitude reports does not run deep: objectual attitude reports share their underlying structure with explicitly propositional reports. The idea is that "Molly wants a burrito" can be analyzed along the lines of "Molly wants that…", where the ellipsis is to be filled in by the description of the state of affairs that would satisfy Molly's desire. If that's right, the familiar treatment of intensionality outlined in the last section can be straightforwardly applied to the case of perceptual reports, (explicitly) propositional and otherwise.

Here are some concerns about the propositionalist approach as applied to perceptual reports. First, following a perceptual verb with a that-clause gives the sentence or clause an epistemic ring.[19] Compare:

(3) He saw the bird flying.
(4) He saw that the bird was flying.

(4) implies recognition of the state of affairs described *as* described, and not merely awareness of it under some guise or other; but not all objectual perceptual reports have this implication. (3) is more easily read extensionally, that is, as compatible with various ways that the subject might interpret the occurrence he is said to see.[20] For instance, we can follow (3) with: "but mistook it for a plastic bag carried by the wind."

---

[17] For more on objectual reports, and for reasons to think these form a distinct category, see e.g. Forbes (2000), Grzankowski (2016).

[18] Quine (1956, 1960) makes the first such suggestion, concerning verbs of search. For recent defenses of propositionalism about intensional verbs, see Liefke (2022) among others. Opponents of propositionalism tend to grant that some apparently objectual attitude reports are correctly analyzed in the way that the propositionalist recommends (e.g., desire ascriptions), but deny that this treatment is appropriate across the board.

[19] Here I am already setting aside purely epistemic uses of perception verbs, whose association with a sensory modality is at most metaphorical (e.g., having been told information that confirms *p*: "I see now that you're right, it's true that *p*"). On the distinction between merely perceptual and perceptual-epistemic uses of the verbs, see Dretske (1969), Barwise (1981), and French (2013).

[20] Note that this pairing is difficult to set up with other verbs of perception. For instance: "he heard that the baby was crying" tends to read as indicating that the subject's information was testimonial in origin. This and other propositional perceptual reports have been appropriated, in English, to play the role that evidentials do in some other languages; in this case, almost to the exclusion of the meaning suggested by taking each of the components of the sentence at face value. This has knock-on effects for how we interpret the objectual "he heard the baby crying": it's often used to convey the meaning that properly belongs to the propositional report. (As pointed out to me by M.G.F. Martin in conversation.) See also Higginbotham's (1983) case for a Davidsonian treatment of bare-infinitive perceptual reports, against (Barwise and Perry 1983); as well as Neale's (1988) critique of Higginbotham.

Brogaard (2018) insists, however, that while *some* propositional perceptual reports have an epistemic dimension, others are "purely perceptual" (142). For instance:

> When one of my cats dies, I lay the cat out on a towel so the other cats can see that their buddy is gone. (*ibid.*)

The rationale behind Brogaard's examples is that they cannot be ascribing knowledge and yet strike us as true; therefore, the sentences have no epistemic import. This conclusion presupposes that knowledge ascriptions are always concept-invoking (cats cannot come to know *that their buddy is gone*), but Brogaard gives no reason for us to think so. Suppose that someone witnesses the theft of a strange-looking object that he himself cannot identify, and I say "He knows that the spectrometer was stolen." If my purpose is to inform the thief that the police might soon be on her tail, this may be a perfectly good report—and one that does ascribe knowledge to the witness. Brogaard's assumption also predicts that sentences like

> The birds went silent; they knew that a storm was approaching.

or

> Young babies know by smell which person is the one that gestated them.

are infelicitous, but that just isn't true. Sometimes what guides our choice of language is not the subject's conceptual apparatus, but that of the speaker and audience.

This does mean that the above point of contrast between propositional and objectual perceptual reports—that the former takes a stand on the subject's conception, whereas the latter needn't—may not be universal. In any case, there is another, more robust difference between the two kinds of reports. Propositional perceptual reports read as factive, and objectual reports do not. One cannot *see that* the baby is crying if there is no baby crying. If that's right, then "Macbeth saw a dagger," used to report a hallucination, is not plausibly equivalent to "Macbeth saw that there was a dagger (before him)."

We might then turn for inspiration to the main rival of propositionalism, namely Montagovian approaches. The latter take their cue from Richard Montague's (1973) suggestion that transitive verbs be given an input of an unexpected syntactic type, one that in effect represents a property of properties.[21] Here we can consider, by way of illustration, Graeme Forbes's (2006) semantics for intensional transitive verbs (ITVs), which draws on Montague's idea among others. If indeed the verbs of perception create intensional contexts, then they might profitably be subjected to Forbes's analysis.

Take the sentence "Gemma seeks a dragon," which displays all three marks of intensionality: failure of existential generalization, substitution-failure, and the

---

[21] Bourget (2017) outlines both a propositionalist and a Montogovian implementation of his claim that perceptual reports are sometimes intensional. He doesn't adjudicate between these; his real concern is the thesis that both approaches would share, namely that the two readings of perceptual reports under discussion (as existentially committed, as existence-neutral) are the product of a scope ambiguity.

specific/non-specific ambiguity. Building on Donald Davidson's (1967) event semantics for action verbs and Terence Parsons' (1990) introduction of semantic θ-roles (which made an earlier appearance in §1.3.1), Forbes suggests the following analysis of the objectual report's specific reading:

> For some event *e*, *e* was a search and Gemma was agent of *e* and a dragon was theme of *e*.

A further move is required to accommodate the non-specific reading, where there is no particular dragon to be the individual in the "theme" role. Here Forbes takes inspiration from Nelson Goodman's (1949, 1976) work on depiction, a context where we also find a version of the specific/non-specific ambiguity: what we call 'a picture of a house' may or may not depict a particular house. When it doesn't, Goodman suggests, the description functions to classify the picture as being of a certain type (a house-picture).

It's not immediately clear, in the case of ITVs, what is being classified. But the event-semantics framework supplies an answer to that question:

> To say that Perseus seeks a gorgon, but no particular one, would, in Goodman's terms, be to classify the event of search as a gorgon-search. To say that Walter needs a bigger boat is to classify his state of need as one of bigger-boat-need. And so on. (78)

This yields an analysis of the non-specific reading of our objectual attitude report:

> For some event *e*, *e* is a dragon-search and Gemma is agent of *e*.

This approach might be serviceable as applied, for example, to visualisations, which can be specific or non-specific in their content (a topic of more extensive discussion in §5.7). But does Forbes' apparatus shed any light on the dual use of perceptual reports for ascribing veridical perceptions and non-veridical experiences? Certainly one can produce interpretations of objectual perceptual reports along the same lines:

> *Specific:* For some event *e*, *e* is a seeing and *S* is subject of *e* and an F is theme of *e*.

> *Non-specific:* For some event *e*, *e* is an F-seeing and S is subject of *e*.

The trouble is that this analysis is appropriate only on the assumption that seeing is the kind of event that can take place with or without an object's presence to the senses. We can make sense of what it is to be seeking a dragon even though there is no particular dragon one is seeking (or, for that matter, no dragons at all). We cannot equally make sense of what it is to be buttering toast when there is no toast to be buttered.[22] On which side does "seeing" fall? In this important respect, we are back at square one.

---

[22] To forestall an objection: one of course can act *as if* buttering toast, whether because engaged in pretence or suffering a delusion. There's still a contrast here. Seeking, unlike toast-buttering, can be non-imitative in the absence of its object. That dragons don't exist is bad news for

The fact that some verbs create intensional contexts while others don't is a reflection of how we conceive of the various states or actions that they name. For instance: verbs of creation are intensional in their progressive aspect only.[23] It can be true that Sara is building a house, without its being true of some house that she is building it. If we say that she built a house, on the other hand, we affirm that there is a house that she built. Building is a way of bringing things into being; one has already begun to build a house when one removes the first shovel-full of dirt from the area where one intends to lay a foundation. But this is an act from which no house need result.[24]

Now consider the following contrast. Our bodily interactions with the world are causal transactions which, if they happen at all, involve us with constituents of reality. They do so whether we know it or not, however we conceive of what we are doing, and so on. This is reflected in our language and our inference patterns. Verbs like to kick (something), to step on, to be hit by, to eat (something) are paradigmatically extensional. On the other hand, attitude reports are paradigmatically intensional, a consequence of the inten*t*ionality of psychological states.[25]

The relationship of perception to these two categories—the psychological, the physical—is ambiguous. Veridical perceptions by definition involve relations to things in the perceiver's environment, whether we spell this out in terms of presentation to the senses or an appropriate causal connection between a representation and what it represents. Yet even in these cases a kinship with psychological states cannot be denied. Perception is conscious awareness; it has a claim on being called mental, even representational, that kicking and eating lack. Our sensitivity to this feature explains the temptation to assimilate perceptual states to psychological states like belief or desire, and the verbs of perception to the psychological verbs. The difficulties with Brogaard's and Bourget's proposals suggest, however, that this might distort the actual workings of our language.

5.3.2 Pragmatic accounts

Let's turn briefly to the other side of the contemporary debate, away from views that intervene on the literal meaning of the verbs of perception. It is a central commitment of these pragmatic accounts, as it was for Austin, that in using perceptual reports to describe hallucinations we make a statement that is not strictly true. There are many ways of building out this scaffolding. I'll introduce two examples.

Gemma's ambitions, but it doesn't mean that she is engaged in *pretend* seeking (its own kind of thing) or that she falsely believes herself to be seeking. The broader point is that we conceive of some states and actions as dependent on their objects in a way that other states and actions are not; and the task is to figure out where seeing stands with regard to this feature.

[23] (Forbes 2006, 37).

[24] The lack of completion entailment for telic verbs in the progressive has been much remarked upon; see discussion of what Dowty (1979) labels the "imperfective paradox," which goes back as far as Aristotle's *Metaphysics*.

[25] On the relationship between inten*s*ionality and inten*t*ionality, see (Anscombe 1965) and (Mackie 1977).

The first is a position that Bourget (2017) reconstructs for criticism, calling it the "unstated qualification view" (UQV).[26] According to the UQV, what one conveys with "Macbeth sees a dagger" is an alternative proposition with the force of *It seems to Macbeth that he sees a dagger*, or *It appears to Macbeth that he sees a dagger*. Of course, the verbs "to seem" or "to appear" themselves create intensional contexts; so the verb of perception can retain its existential import without the whole sentence being existentially committed. The unstated qualification view is then a candidate explanation, consistent with the claim that the verbs of perception create extensional contexts, of how perceptual reports can behave as though they were *in*tensional.

The unstated qualification view involves covert additional content that alters the meaning of a perceptual report. Daniel Hoek (2018) appeals instead to a type of pragmatic weakening, "conversational exculpature," in which some of the usual consequences of a statement are dropped (152). A reader of Shakespeare might say:

> The dagger Macbeth saw in front of him was covered in blood stains. (179)

The reader then speaks "as if the vision were veridical" (*ibid*.), though recognizing that it isn't. Hoek presents this as an instance of a general tendency to "speak as if any representation we are describing is accurate, whether or not we believe this" (*ibid*.). That supposition of accuracy is subtracted in what Hoek labels a *veracity exculpature*. What gets communicated is then "a message […] that concerns only the content of Macbeth's visual experience, and therefore fails to entail that there was a dagger in front of Macbeth" (180).

As Hoek recognizes, it's difficult to say exactly what that message is. He makes two attempts. First:

> Macbeth's visual experience was as it would be if he were to see a bloody dagger. (181)

Or, avoiding the counterfactual:

> Macbeth's visual phenomenology depicts a bloody dagger. (182)

Hoek also considers that "exculpature can expand our expressive range" and that there may be "no natural literal form of expression that precisely captures" what we convey in asserting that Macbeth saw a bloody dagger (*ibid*.).

What I want to highlight about these proposals, for the moment, is the nature of the contrast between them and the semantic accounts we considered in the previous subsection. (I reserve specific criticisms for a later point in the chapter, §5.8). Both seek out a proposition that the second, more puzzling use of perceptual reports communicates. Parties on both sides of the debate could in principle agree on exactly what content is communicated by "Macbeth sees a dagger" on the usage at issue; notice that Hoek appears, like Brogaard and Bourget, to have the intentionalist's account of the nature of perception front of mind (as evidenced by his talk of the content of experience, and the

---

[26] Bourget cites Dretske (1969), Grice (1989), and Soames (2003) as its main adherents. Of these authors, however, only Soames unambiguously espouses the view that Bourget describes.

idea that visual phenomenology is *depictive*, i.e., representational in character). The clear point of disagreement concerns *how* the content is communicated.

5.4 A diagnosis

Taking in the development of the debate about the perceptual verbs from a remove, I think we have good reason to suspect that it has run up against a dead end. First, there is the extent to which the early debate has simply replayed itself against a new philosophical background. Second, one can imagine the back-and-forth between those who want to take both uses of perceptual reports literally and those who characterize one use as non-literal continuing indefinitely into the future, as proposals on both sides are refined in response to counterexamples. Third, the contested question is *not that interesting*, at least from the point of view of a philosopher of mind; odd when we consider how much attention it has attracted from precisely that demographic. I suggest that this is because the existing debate acquires an illusion of significance from its status as a proxy war for disagreements about the metaphysics of sensory experience. What is going on here?

I contend that the root problem with the debate about the perceptual verbs is that the framework imposed on it by the sense-datum theorists has never really been thrown off. Looking back at Moore's assessment of our "sees" talk, we can notice that he characterizes it as ranging over two domains: subjective appearance (constituted by sense-data) and objective reality (constituted by physical things). The independence of the former from the latter is supposed to be revealed by the possibility of illusion and hallucination, states that we may be unable to distinguish from veridical perceptions on a purely introspective basis. This metaphysical picture sets up a certain expectation about our expressive powers: we should have no trouble providing an independent description of the subject's point of view, no matter how disconnected her experience may be from the actual state of her environment.

Neither the general structure of Moore's account nor the expectation it gives rise to is unique to the sense-datum theorists. Both recur in the intentionalist theories of perception now dominant, which of course posit their own "common factor" between veridical perceptions and hallucinations.[1] Both can be traced back to Descartes who, as we saw in §2.1.2.1, persists in affirming that *it seems to him* that he sees and hears such-and-such objects even when he has cast doubt on the reliability of perception as a source of knowledge about the external world. "That," he assures us, "cannot be false", and is what "is properly called sensation" (AT VII:28/HR 1:153, trans. Williams 1978, 63) These distinctively *philosophical* ideas have been allowed to guide the interpretation of our everyday linguistic behaviour, with unfortunate results.

Moore, for his part, is explicit about treating the "merely sensory" as a distinct tract of reality. But when Ayer claims, in *The Foundations of Empirical Knowledge*, that there are two senses of "see," he purports to have done away with any commitment to sense-data: the second sense of "see" is one on which "it is not possible that anything

[1] The representational content of a sensory state is what is supposed to determine the subject's phenomenology, and so it presents itself, much as an array of sense-data would, as the obvious target of descriptions of experience: when we say that Macbeth sees a dagger, the thinking goes, we must be talking about *that*.

should seem to have qualities that it does not really have, but *also not necessary that what is seen should really exist*" (23; my emphasis). The upshot is that one can speak truly, in saying what one sees, even if what one claims to see is not present in the environment. Notice, however, that this is is just the sense-datum theorist's claim in new packaging: the ontological commitment has been removed by fiat, with nothing to take its place.

As Austin emphasizes, the fact that Ayer can sincerely diagnose the sense-datum theorists as making a *linguistic* proposal confirms that Ayer's own view is a terminological variant on that much-maligned account of the metaphysics of perception. Yet when Austin rails against the "two senses" hypothesis, he does not question his opponent's assumptions about what we *mean* when we say that Macbeth sees a dagger. As I explained above, Austin defends the integrity of the perceptual verbs simply by pushing this supposed other message into the sphere of the non-literal, and by downplaying the frequency and the seriousness with which we make use of perceptual reports in this non-literal way.

These early moves established a dialectic that persists to the present day, one whose defining question is whether to assign a certain proposition to the semantics or the pragmatics of perceptual reports. Against this background, innovation takes one of three forms: a proposal about what exactly that proposition is, an account of the semantics of perceptual reports that explains how the proposition is expressed, or a suggestion about the pragmatic mechanism that gets the proposition across. The fact that none of these refinements has been entirely satisfactory is one indication that the literature might rest on a mistake. But it will take more doing to show that *no* account in the usual style can free us of a lingering puzzlement about the meaning of "see."

5.5 Perceptual reports are not ambiguous

Our first task is to get clear about the literal meaning of the verbs of perception. While the standard accounts ultimately commit the same error, they do so in different ways: the hypothesis that perceptual reports are ambiguous makes a false prediction about the workings of the verbs of perception that a pragmatic account need not. This is so whether the proposed ambiguity is lexical, as the likes of Moore and Ayer assumed, or scopal, as Brogaard and Bourget have recently argued: despite some interesting differences between these views, they are vulnerable to the same fundamental objection, and I group them together on that basis.

Recall that attitude verbs generate an ambiguity between *specific* and *non-specific* readings of indefinite descriptions and of certain quantifiers that fall in their scope.[2] "He wants a pet unicorn" can attribute a desire for a particular unicorn, or merely the wish to take possession of some unicorn or other. The former is the specific, and the latter the non-specific, reading. In contrast, the extensional context of "He owns a unicorn" forces a specific interpretation. It does so because *owning* is unambiguously a relation between one thing and another.

---

[2] Again, in Quine's terminology, these are "relational" (specific) and "notional" (non-specific) (1956, 177).

Suppose that Macbeth is really seeing—veridically perceiving—a dagger. Can it be that he isn't seeing any dagger in particular? No: the perceptual relation determines one entity as relevant. Compare, next, the case where Macbeth is hallucinating and this is common knowledge between us. I say, "Macbeth sees a dagger." It's uncontroversial that my utterance may be apt even though Macbeth's experience neither depicts nor is the causal product of any particular dagger. In this way the case is analogous to that of visualizing a dagger or desiring one. But whereas one can *desire*, e.g., "mere relief from slooplessness" (Quine 1956, 179), there is no *seeing something or other*, in the relevant sense of that phrase.

I intend, for the moment, simply to be making a linguistic observation: our language does not provide for the possibility of an experiential analogue of non-specific desire. If I say, "I need to eat something," the most salient interpretation of that sentence is that I have a non-specific need; one that can be satisfied by my going into the kitchen and grabbing the food item I encounter first, whichever it may be. If I add that I am eating something or other, or answer the question, "What are you eating?" with, "Nothing in particular," it's understood that I simply *have not specified* which thing I am eating. Perhaps I am unwisely consuming an unidentified object, or I can't be bothered to elaborate. But if I really am eating, there inevitably is some particular, some specific thing, which is the recipient of my action.

The verbs of perception have the same effect as the verb "to eat" used transitively: attempts to ascribe a non-specific hallucination wind up conveying either ignorance or deliberate vagueness about the nature of what is experienced. The inevitability of a specific reading of the object of a verb, even in the face of ignorance on the subject's part about the nature or identity of what she is involved with, is a hallmark of causal relations. It's because perception is a way of being affected by things in one's environment that descriptions of *veridical* perceptions have this feature. The fact that it persists even when the intention is to ascribe a hallucination—when there is patently nothing that the subject perceives—suggests that the speaker is, at the level of literal meaning, still making a relational claim. She may yet be using that claim for a novel purpose, e.g., that of saying p in order to convey q, a possibility I will target in the sections that follow. But the hypothesis that perceptual reports between relational and non-relational readings can be ruled out.

5.6 The interdependence of the two uses

Contra Austin, it cannot reasonably be denied that there is a fully legitimate and familiar use of perceptual reports for characterizing hallucinations, illusions, and mere appearances generally. That being so, the crucial question is what such a use *requires*. A natural thought is that it calls for an interpretation of perceptual reports, whether literal or non-literal, that does not imply the existence of whatever the object phrase purports to single out. If we just do away with this business about a perceptual relation, the thinking goes, we will be left with a suitable description of the hallucinating subject's experience.

Despite its intuitive appeal, this strategy misconstrues the way that perceptual reports fit into the project of describing the subjective point of view. The existential commitment standardly associated with the use of a perceptual verb is not a feature independent of the role these verbs play in evoking a subject's phenomenology. Instead,

claims about what the subject is related to *themselves* function as descriptions of sensory experience.

This is a neglected insight of Anscombe's (1965) paper on the language of perception, and one that comes to the fore in the following passage:

> Suppose a bright red plastic toy elephant looks greyish brown to me in a certain light. Only if I do not know that the greyish brown colour is mere appearance do I say without any special context (*e.g.* that of describing impressions), or apology, or humour: 'I see a greyish brown plastic toy elephant.' This is because we understand the description-of-an-appearance 'greyish brown' by understanding the description 'greyish brown': this describes what the appearance is *of*. To do that, it must in the first instance be a description of such a thing as it would be true of (for the appearance is an appearance of that)—really, and not merely in appearance: this will be its primary application. (172)

The key observation here is that "I see a greyish brown toy elephant," as spoken "without any special context […], apology, or humour," is neither a comment merely on the properties of the toy one sees, nor merely on one's visual experience. The remark makes no distinction between these; and that stands to reason, since in cases of veridical perception no such distinction is generally assumed. (Another aspect of the passage is a claim about fundamentality: greyish brown is *in the first instance* a description of greyish brown things. The appearance-characterizing use is secondary, parasitic on, this "primary use.")[3]

As Anscombe's student, Hidé Ishiguro, emphasizes: "in normal conditions when I say 'I see *X*' I am making a claim both about the presence of *X* […] and about my mental state or visual experience" (1967: 54-5). If it's by making a claim about the presence of *X* that one describes what a veridical perception is like from the subject's perspective—if the description of the perceptual relation *constitutes* a description of the phenomenology—then the strategy of casting that claim aside to generate a description of mere sensory appearances is ill-conceived. In particular, we should expect any change to the content that follows the perceptual verb to make a difference, not just to which objective state of affairs is thereby reported (successful perception, mere hallucination), but also to how the experience is characterized as being from the first-person point of view.[4] While the first of these changes is desirable, the second is not; for the relationship between perception and hallucination is not just one of objective difference, but also one of subjective resemblance.

Perceptual verbs already belong, in their most ordinary uses, to the project

---

[3] The target here is sense-datum theory but especially phenomenalism. Wilfrid Sellars (1956) makes a similar point in his own discussion of "looks"-statements: the notion of *looking F* is dependent on the notion of *being F*, and not, as the phenomenalists would have us believe, vice-versa. He claims that "to say that 'X looks green to S at t' is, in effect, to say that S has that kind of experience which, if one were prepared to endorse the propositional claim it involves, one would characterize as *seeing x to be green at t*" and "when I say 'X looks green to me now' I am *reporting* the fact that my experience is, so to speak, intrinsically, *as an experience*, indistinguishable from a veridical one of seeing that x is green" (273).

[4] An idea that Clarke (1965) introduces in his discussion of the claim, on the part of some sense-datum theorists, that we only see the facing surfaces of physical objects.

of characterizing the conscious states of subjects. To describe an experience, we cite the objects and sensible qualities that (actually or habitually) give rise to it, and the perceptual modality through which they are encountered. We persist in doing just this whether or not those objects are present, or those qualities instantiated, in the subject's environment. That's not an indication that perceptual reports have two meanings, or that the presentation to the senses of an object is merely an optional accompaniment of an instance of seeing. It's rather a sign that there is only one tool available for the job we want to perform.

5.7 Phenomenological particularity

I intend the moral of the last section to be quite general. It can nonetheless be sharpened by reference to one aspect of the subjective character of sensory states, which has a close connection to the issue of specificity discussed above: their *phenomenological particularity*.[5] I'll bring this feature into view by recapitulating the salient points of an argument that Martin (2020) develops to a slightly different end.[6] The argument exploits a comparison between sensory experience (perceiving, hallucinating) and sensory imagination.

First consider a case of visualizing an invented sculpture, the likes of which has never existed (732). We can observe that an empty singular term would be a useful device for specifying the content of this experience.[7] On the one hand, we want more than a merely existential condition. One can *imagine that there is a sculpture with such-and-such properties* without having a sculpture before the mind's eye, for instance by visualizing a cardboard box in which the sculpture has been packaged for transport. On the other hand, we're supposing that the sculpture doesn't exist, so it's not as though we're in a position genuinely to refer to it.

An empty singular term strikes this compromise, capturing the "presence in absence" characteristic of imagining (Sartre 1940, 229). It does so, moreover, without bringing in a further, unwanted element: the implication that there is a *particular* statue that one visualizes. While there is such a thing as visualizing a particular statue, real or invented, there is also such a thing as merely visualizing some sculpture or other. This latter possibility manifests in the fact that there are cases where it just doesn't make sense to insist on an answer to the question *which statue* is being imagined—cases where such insistence would reveal a confusion, on the part of the inquirer, about the nature of imagination.

If I am perceiving a sculpture, on the other hand, it will be obvious to me that the corresponding "which one" question has an answer, whether or not I have much to offer

---

[5] Among those who emphasize the particularity of perception from a phenomenological perspective are Soteriou (2000), Martin (2002b), Montague (2011), and Schellenberg (2018).

[6] Martin's target is Schellenberg's (2018) account of phenomenological particularity. In that context, the point of the argument is to show that her proposed condition for the manifestation of that feature—that it "phenomenally seem to the subject that there is a particular present" (17)—is too weak.

[7] Compare the "potentially gappy *de re* modes of presentation" that Schellenberg appeals to in her account of perceptual content (2010, 34–43), as well as the use that Sainsbury (2005) makes of empty singular terms.

in the way of an informative response.[8] It will be obvious not just conceptually, in that I understand what the ingredients of successful perception are, but phenomenologically: the object of a perception is *right there* for the subject, determinately located in space and time, in a way that the object of a visualization manifestly is not. This contrast between visualization and perception matters because it is a critical source of leverage against the widespread understanding of sensory phenomenology as exhausted by the ways that experiences can qualitatively resemble and differ from each other—and because that conception, as will continue to emerge, is what prevents us from grasping the significance of statements like "Macbeth sees a dagger." [9] I'll consequently reinforce the claim at issue before continuing with Martin's argument.

That there is an introspectable difference between visualization and perception is, I trust, evident enough. But one might be inclined, perhaps recalling Hume's talk of "force and vivacity" ([1739] 1965; T 1.1.1.3), to push back on the portrayal just offered of that difference. Is this not a difference in *intensity*, which is itself a qualitative matter? This isn't the objection that it first seems to be, as the following may help to bring out. People with hyperphantasia describe having some imagery that is "as vivid as real seeing" (Zeman et al. 2020), and there are anecdotal reports of such imaginings being mistaken for perceptions.[10] (The veracity of these reports does not really matter for my purposes; I merely use them to inspire a thought experiment.) What is the nature of the error that these subjects describe? It's not one that *originates* as a false judgment about the type of mental state they are in. Rather, they claim to experience the presence of something in their environment—an object seen, a sound heard—when in fact, as they subsequently realize, there is nothing there. (It's upon making such a realization that one reconsiders the presumptive status of one's experience as genuinely perceptual.) You can notice for yourself that imagining a visualization (of object *O*) with the vividness and determinacy of a perception yields the very same result as imagining that *O* there in front of you, occupying a position in your visual field. Even if perception and visualization do each occupy positions on a spectrum of vicaciousness or force, the fact remains that an experience with the vividness of a perception will have the character, for the subject, of an encounter with something in the environment—not the character that typical (non-hyperphantic) visualizations have, which is transparently that of a representation.

Let's now put to work our observation about the introspectively evident contrast between such visualizations and perceptions. The phenomenological particularity of successful perceptions is a feature that hallucinations inherit, as evidenced by the fact that a subject who *doesn't know* she is hallucinating will also feel the "which one?" question to be appropriate (Martin 2020, 732). As far she can tell, she is in a position to answer it, if only by demonstrative reference. That's to say that whereas the phenomenology of a typical visualization advertises its own disconnection from reality, the phenomenology of a hallucination dissimulates it: as Martin has elsewhere emphasized, these experiences are "misleading as to how the subject's environment is, but also […] misleading about themselves" (2002b, 401). One therefore has to go beyond the subjective character of a

---

[8] "I don't know which statue it was, I'm terrible at recognizing artworks"—remarks like these reflect a failure to *identify* the particular at issue rather than a failure to appreciate that a particular *is* at issue.

[9] Compare (Martin 2002a, 186).

[10] See, e.g., (MacFarquhar 2025).

hallucination—drawing on other sensory modalities, considering what one knows about ways that the world can be, asking someone else what they see—to diagnose one's experience as the counterfeit perception that it is, and to recognize that the "which one" question may not have an answer after all.

The phenomenological contrast that Martin emphasizes has implications for what it takes to faithfully describe the various states (visualization, perception, hallucination) at issue. We determined that it was appropriate to use an empty singular term to specify the content of certain visualizations. Now, one might have thought that empty terms would also be well suited for reports of hallucination, since these experiences are not actually perceptual encounters with anything. But introspective reflection on the difference between what it is like to visualize a sculpture and what it is like to perceive or hallucinate a sculpture shows that what goes for sensory imagining does *not* go for sensory experience. To capture what is distinctive about the latter, to describe perception and hallucination in a way that differentiates their character from that of sensory imagination, we need more than what a potentially empty term can provide: we need singular *reference*.

Here's why. Singular reference—actual or purported—commits the speaker to the particularity of the sensory state being described because it's a way of identifying a particular thing *as* the object of the experience. When you say in all sincerity that you want *that burrito*, you are making it clear that you take your desire to be directed towards an individual and not just any member of a given kind; you make this clear even if your attempt at demonstrative reference turns out, for whatever reason, not to go through. If we allow that there can be terms which are both singular and empty, however, the fact that a singular term has been invoked no longer has that implication; the link between singularity and particulars is broken (Martin 2020, 732–3). If this were all we had to work with, we would be left unable to capture the distinctive way that particulars figure in the phenomenology of perception and hallucination alike.

It's uncontroversial that the governing aim of the use of perceptual reports under discussion is to evoke the subjective character of hallucinatory experiences. We capture the phenomenological particularity of veridical perceptions by referring to an individual as the object of the perceptual relation. If indeed "Macbeth sees a dagger" is a good way to evoke his fatal vision, as it very much seems to be, it's because it has the same implications as any other perceptual report. I'll turn to the problem that this introduces in §5.10. But first I need to block off a tempting avenue of escape.

5.8 Candidate propositions

Up until now I have left it open that perceptual reports might be used, in a way that exploits their literal meaning, to convey a truth about the phenomenology of a hallucination. At this stage we have what we need to identify the faults in this second, more flexible proposal, and so to abandon once and for all the shared assumption I targeted in the introduction.

In the last section, I argued that particulars need to come into our descriptions of perceptions and hallucinations alike. Yet one might think it remains to be determined exactly how. A solution that immediately presents itself is to take an ordinary perceptual report and embed it under the scope of an intensional operator. The various attempts that

have been made to spell out what we really mean when we say that Macbeth sees a dagger are different ways of enacting this solution. Perhaps what we get across is a claim about how things appear to the subject ("It seems to Macbeth that he sees a dagger"), or a message that invokes a counterfactual situation ("Macbeth's experience is as it would be if he were seeing a dagger"), or a statement that exploits notions of representation or depiction ("Macbeth sees an array of sense-data such as are characteristically produced by the presence of a dagger," "Macbeth's experience represents to him that there is a dagger present").[11]

In each case, the underlying strategy is the same. The inclusion of a perceptual verb, as I emphasized in §5.5, forces a specific reading of the indefinite description. So an advocate for one of these paraphrases can assure us that when we say that it seems to Macbeth that he sees a dagger—I'll use the "seemings" proposal as a representative for brevity's sake—the situation that we describe him as seeming to be in is one that features a particular dagger. At the same time, the embedding means that we get to say something true; and we get this result at, apparently, no cost.

Particulars are indeed being "brought in" here, but not in the requisite way. First, notice that a subject who describes their experience as one of *seeming to see a dagger*, like the subject who merely visualizes a dagger, can reject as inappropriate the question, "Which dagger?". The subject who sincerely claims to perceive a dagger cannot, and the hallucinator is in the same position *subjectively speaking*.

This makes it vivid that in merely comparing a given experience, *E*, to an instance of seeing, we are not characterizing *E* itself as a confrontation with a particular. We are saying that it is similar, in some interesting respect, to a case of seeing something. If the comparison is apt, this respect must be one common to hallucinating a dagger and seeing a dagger, or between *this* seeing of a dagger and *that* seeing of a (quantitatively distinct) dagger. The description we are providing in this comparative way, then, is merely qualitative. Yet one of the lessons of §5.7 was that a characterization of perceptual (and hence hallucinatory) phenomenology in purely qualitative terms is not satisfactory.

If we want to know what it is like to hallucinate, our best guide is the person who is hallucinating but doesn't realize it: the subject who, having no reason to think that her situation is unusual, feels no pressure to bring her report into conformity with the objective facts of the situation. This hallucinator will behave exactly as the perceiver does. She will make a claim about the presence of things to her senses, not a claim about the subjective resemblance of her current situation to a situation in which she confronts a particular. Issuing as it does from a first-personal grasp on her own experience, this choice of words might signal something important.

To be clear, I affirm that claims like "It seems to Macbeth that he sees a dagger" can be true in circumstances of hallucination; I also affirm that ordinary speakers sometimes make such claims. What I deny is that these statements are interchangeable with "Macbeth sees a dagger." I deny that they are suitable as reports of the experience that "Macbeth sees a dagger" so successfully evokes. They are accurate descriptions of *a* subjective situation, but not of the right one. They serve *a* purpose, but not the same

---

[11] The first of these paraphrases is articulated most explicitly by Soames (2003, 184) and is arguably hinted at by Grice (1989, 44–8); it comes in for criticism from Bourget under the heading of the "unstated qualification view" (2017, 520). Hoek considers the second and a version of the third (2018, 181–182).

purpose as the kinds of remarks that set the debate about the verbs of perception in motion.

5.9 An innocent adjustment?

The idea I have been picking on is that while "I see an *F*" makes a claim about the way that I am embedded in my environment which "I seem to see an *F*" does not, these two claims are nonetheless reports of the same subjective state. This is a tenacious view, and I now want to consider—and reject—a further line of response that might be pursued in defence of it.

P.F. Strawson (1979) imagines some philosophers who, happening upon an innocent victim, ask him to describe his visual experience. The "non-philosophical observer" (henceforth "NP") replies: "I see the dappled deer grazing in groups on the vivid green grass" (*ibid.*). Unsatisfied with this naïve response, NP's interrogators continue:

> We explain that we want [NP] to amend his account so that, *without any sacrifice of fidelity to the experience as actually enjoyed*, it nevertheless sheds all that heavy load of commitment to propositions about the world which was carried by the description he gave. We want an account which *confines itself strictly within the limits of the subjective episode*, an account which would remain true even if he had seen nothing of what he claimed to see, even if he had been subject to total illusion. (127–8; my emphasis)

The impulse I am targeting is on full display here: to concern oneself with the subjective fact is, one supposes, to move away from talk about deer and grass.

As philosophers, we are accustomed to certain ways of responding to a request of the kind that Strawson outlines. There is the description in terms of "lights and colours, patches and patterns" favoured by the sense-datum theorists (128), the claim about how one's experience represents the environment as being, the hesitant "It seems to me that I see…" But NP takes another approach, which I join Strawson in endorsing:

> I had a visual experience such as it would have been natural to describe by saying that I saw, etc. . . . were it not for the obligation to exclude commitment to propositions about independently existing objects. (*ibid.*)

NP realizes that the best way to answer the sceptic's question—to answer it, as the sceptic's own instructions require, without doing violence to the character of the experience at issue—is to set aside the possibilities that that question itself introduces, not incorporate them into the report's content.

Defenders of the "seems"-paraphrase will concede that of course there is a difference between reporting one's experience as a case of *seeing the deer* and reporting it as a case of *seeming to see the deer*. Yet they will insist that this change has no implications for the phenomenology of the state that one reports. The only difference here is *epistemic*. A subject commits herself to less in making the second statement than in making the first; she is simply hedging her bets.

This rejoinder misses what is so distinctive about sensory experience as compared to our other mental states, and what allows perception to play the decisive role that it

does in our judgements about the world. As Austin emphasized, perceiving a thing does not normally have the status of evidence for that thing's presence (1962, 113). Noticing some hoofprints is a way of acquiring evidence as to the presence of a pig. But when you see the *pig*—well, there it is, case closed. Perception has the effect of convincing us, unless we have been given independent reason to doubt whether really are perceiving, that the pig is before us. It does so not, or at least not only, because of a background hypothesis about the reliability of this faculty; it's the subjective character of perception itself, the "what it's like," that has this effect.

This observation introduces a requirement on a correct account of sensory phenomenology: we need a description of the phenomenal character that predicts its actual impact on our doxastic state. If what we called "seeing" were an experience that only suggested the presence a thing, one's seeing of the pig would at best add more evidence to the pile. Some of our experiences *are* like this, whether because of problems with our sense organs (myopia, partial deafness) or because of environmental conditions (low light, fog, an unusual confluence of tastes, background noise). And yet most perceptions do not have the character of a mere suggestion. Here one might double down: doesn't the possibility of mistaking a hallucination for a perception just establish, supposed phenomenological subtleties aside, that *seeings* are fundamentally *seeming seeings*? No. That conclusion mistakes the outcome of reflection on ways that we can be mistaken for some kind of tentativeness or equivocality in the character of the original state.

If we treat *seeming to see (something)* as the basic, properly "mental" state, and think of perception as the combination of that state with a fortuitous external circumstance, there is no making sense of Austin's observation. Seeming to see is a state that, while it invites us to judge that things are one way, is patently compatible with their being otherwise. But what manifests itself to you as mere (i.e., inconclusive) evidence for the truth of a proposition does not rationalize the conclusion that the proposition is true. Accounting for the impact of perception and hallucination on our beliefs—notice that my immediate concern is with psychological dynamics, not the actual epistemic standing of the various states at issue—requires a different picture than the "seemings-first" conception provides.

The best description of what it is like to perceive, as I have been arguing, is a description in terms of the subject's awareness of what is around her: "Ah, here's the ring I thought I lost. Oh no, there's that irritating person I've been trying to avoid." These are the kinds of things that veridical perceptions, and hallucinations or illusions that we mistake for veridical perceptions, regularly induce us to say. A picture on which ordinary things themselves are constituents of our experiences when we perceive, which treats hallucinations as entirely parasitic on perception so understood, predicts that both states would settle questions for us in the way Austin describes.

But then—returning to the suggestion that set off the present line of argument—the epistemic import of the fact that I am seeing a certain object is not something separate that can simply factored out. Our description of the phenomenology already includes the presence of the object. A description that leaves it out gets to be true across a wider range of circumstances, but the experience one thereby ascribes is an anemic counterpart of the original. We cannot move from "Macbeth sees a dagger" to "It seems to Macbeth that he sees a dagger" without *also* characterizing Macbeth as being, from his own point of view,

in an importantly different mental state than the unmodified report has the potential to ascribe.

5.10 Same claim, different activity

We've now come face-to-face with a fundamental tension. On the one hand, we must account for the fact that the speaker of "Macbeth sees a dagger" neither believes nor wants their audience to believe that there is a dagger before Macbeth. On the other hand, given the role that perceptual reports have in our quest to express and understand the subjective point of view, we seem to be called not just to say, but in a sense to *mean*, the very same thing when we invoke the verb "see" to talk about Macbeth's hallucination as when we invoke it to describe his prior, undeluded perceptions: that's what it takes to get the character of his experience *right*. It's now becoming clear, I hope, that the strategy I've applied elsewhere in this dissertation can also serve us here.

Most of the literature on the verbs of perception is concerned with the question of whether a certain proposition is stated or conveyed by a class of utterances. This is one dimension of variation: how a content is communicated. But the emergence of the field of pragmatics was driven in part by an increased sensitivity to the kind of action that an utterance constitutes, or with what speakers are trying to *do* with their words. To make sense of the way that we report hallucinations we need to reconsider, not what the speaker is saying, but to the nature of the speech act they perform.

Drestke's (1969) comparison between talk about hallucinations and talk about dreams should, I think, be interpreted in this spirit. Put in the service of reporting a dream, the sentence "I held a butterfly" plays a different role than it normally does. The clearest indication of this is that it ceases to imply the real-world existence of a butterfly, though crucially it does retain that implication *within* the dream narrative. Dretske observes that by making our purposes known to our audience—whether that's by prefacing a narrative with a signaling phrase ("in my dream," "once upon a time") or by any number of other contextual cues—we can alert them to the fact that what is about to occur is not ordinary speech. I emphasize that the suggestion is *not* that there is a hidden operator at work when we say that Macbeth sees a dagger—another instance of the approach I have been militating against in the preceding sections. The moral is rather that we may be *doing* something quite different than what we do when we report veridical perceptions.

Turning our attention on the uses of perceptual verbs that form the preoccupation of this chapter, we philosophers are like an anthropologist who continues her record of everyday human behaviour without realizing that what she has happened upon is a rehearsal for a play. The actions she observes may be the very same as those characteristically performed in, say, a fight, but their significance differs, because they belong to an activity other than fighting. What first strikes us about remarks like the one Moore extracts from Shakespeare is that the perceptual report is appropriate even though we know that no object (at least, no object meeting the description or answering to the name) is perceived. Freedom from existential commitment where we would otherwise expect it can indeed be a sign of a lexical ambiguity, of intensionality, or of implicature. As emerged in chapter 3, however, it can also indicate that speakers are engaging in *pretence*. In this case, the existential import of a remark is not removed so much as

displaced, from reality to the world of the game. This feature will be the key to solving the problem we face in trying to do justice to the phenomenology of hallucinations, and we can invoke it to explain why we often report these experiences in terms proper to veridical perception.

5.11 Pretence in the sensory context

Recall the key components of Kendall Walton's account that I isolated in chapter 3 (§3.4):

> *(i)* Props combine with principles of generation to generate truths-in-the pretence.
> *(ii)* One may use elements of one's own or of another person's psychology as a prop.
> *(iii)* A pretence may have as its initiating stipulation that some misinformation is veridical.
> *(iv)* The stipulation of *(iii)* makes available, within the pretence, the kinds of thoughts, attitudes, and emotions characteristically produced by encounters (whether first- or second-hand) with things in the world.

I propose to understand the use of perceptual reports that has spawned so much discussion on this model. Consider a subject who, unaware that she is hallucinating, finds herself in a state that she would express by saying that she *sees a persimmon*. The hallucination is, in Evans's terminology, a perceptual informational state (1982: 123). This representation is an element of her psychology compare *(i)*) that represents the world as being a certain way (compare *(ii)*): the basis for our pretence is no longer a work of fiction, but the sort of empty representation produced and sustained by the mind itself. If we pretend that this subject's informational state is veridical (compare *(iii)*), we thereby pretend, without any further stipulations on our part (as in *(iv)*):

- that the subject possesses information deriving from an individual persimmon with such-and-such properties,
- that there are singular propositions concerning this particular fruit,
- that the subject can entertain and express these propositions and that we can do so as well,
- that we can understand the hallucinating subject's remarks,

and so on. In brief, we recreate internal to the pretence the structure characteristic of the good case, a structure that is normally provided by the entity that a singular thought concerns. And this puts us in a position to engage our pre-existing cognitive and linguistic capacities, as, for instance, one can best engage certain soccer-playing capacities by pretending that there is a ball that one is kicking around.

Macbeth has information in the form of a sensory representation that purports to derive from an object. He says, on this basis, "I see a dagger." Now, what he says is just wrong: seeing an object requires the presence of that object to the sense of sight. But his erroneous claim is also evocative. We cannot ourselves assert sincerely that he sees a dagger, because we know too much. But we can get on the same page as Macbeth by *pretending* what he innocently *assumes*, namely that his informational state is veridical.

By using Macbeth's psychological state as a prop, we are able to do something that

we cannot do in any other way: engage with Macbeth's point of view without molding it to fit into our picture of objective reality (to fit, in effect, *our* point of view). "It seems to Macbeth that he sees a dagger" exemplifies such molding-to-fit. If we refuse to pretend, and confine ourselves to the communication of true propositions, the best we can muster is a comment on the subject's mental life from the outside: the kind of comment that we make when we say that the child thinks that there is a magical man who lives at the North Pole. Only within the boundaries of a pretence can we move beyond generalities to think and talk about *the very dagger*, the individual, that Macbeth is (as we also pretend in our efforts to understand him) concerned about.

Again, one *can* say, "It seems to Macbeth that he sees a dagger," or, "Macbeth is having a visual experience as of a dagger," or make any number of similar statements. What one thereby describes are states that are neutral as concerns the presence of a dagger. But from Macbeth's own point of view, he is not in a state that is compatible with the absence of a dagger. From his point of view, his visual experience is in direct competition with the rest of what he knows. The continuation of his soliloquy makes this clear:

> I have thee not, and yet I see thee still. / Art thou not, fatal vision, sensible / To feeling as to sight? Or art thou but / A dagger of the mind, a false creation, / Proceeding from the heat-oppressed brain? / I see thee yet, in form as palpable / As this which I now draw. (Act 2, Scene 1)

Even as Macbeth becomes convinced that he is merely hallucinating, he speaks as though he is seeing a dagger. He asks whether *that dagger, which he sees,* is a product of the heat-oppressed brain. (He cannot plausibly be interpreted as asking, contrary to Moore's telling, whether the sense-datum of which he is immediately aware is the causal product of a dagger or a fever.) Macbeth is resolutely faithful to his visual phenomenology, though the sum total of his knowledge leads him to refrain from making the judgments about the room around him that, in the absence of such further knowledge, his experience would induce him to make. This is what it is to focus in on the subjective perspective: it's to *set aside* certain sceptical possibilities, not adjust one's report to accommodate them.

I've emphasized throughout this chapter the importance of sensory phenomenology in leading us to a right account of what is going on with our uses of the perceptual verbs; I have focused on the way that *particulars* figure in perception and (derivatively) in hallucination alike, and on the power, in this context, of existentially creative make-believe. This was because the qualitative aspect of sensory phenomenology has been emphasized to the exclusion of the element of particularity. But the qualitative character of experiences is also part of what is at issue when we engage with the perspectives of other people, and pretence can be invoked to illuminate this side of things as well.

There is a convergence, in the various attempts at paraphrase that we've considered, towards the idea that the uses of perceptual reports I've classified as non-literal in some way involve a comparison with actual, imagined, or counterfactual veridical perceptions. There is something right about the spirit of these suggestions. To engage in the kind of pretence I have described is to engage in perspective-taking, to use one's own, first-personal grasp of what it is like to encounter the world through the senses as a basis for understanding the experiences of others.

Consider, first, the word "looks" and its counterparts for the other sense modalities,

which are generally agreed to have a comparative use. [12] For instance: "the flowers look fresh" can mean that the flowers look as fresh ones do. There are also explicitly comparative constructions featuring these expressions: "this bit of leaf looks like a musketeer," "Saskatoon berries taste like blueberries, but with a hint of marzipan." These statements help us to expand our understanding of reality using our existing stock of phenomenal concepts.

These comparative claims are just that—claims, about what resembles what—but they also offer an instruction. To imagine what a Saskatoon berry tastes like, you should imagine tasting blueberry and marzipan. This, presumably, is what makes these statements as informative as they sometimes can be. If you're not familiar with the look of a musketeer, or the taste of blueberries and marzipan, merely grasping the proposition expressed by the comparative statements I've just supplied will be of little use to you.

What we call "hallucination" is a dimension of the mental that in some measure escapes our understanding. Here one might feel resistance: aren't hallucinations very easy to describe? Hallucinating is like perceiving, without the object; hallucinating an *F* (visually, let's suppose) is like seeing an *F*, except there's no *F* that you're seeing. Yet what is left to seeing an *F*, once we take away the seeing of an *F*? This gloss is unquestionably tempting, but it's also bizarre—bizarre, that is, if we take it as offering a direct description of what it is to hallucinate. There is another way to understand its function: this 'description' tells us what to imagine seeing, and what not to imagine being the case, in order to get a grip on a particular experience. When the context rules out veridical perception, the statement that *S* sees an *F* can have the same effect.

Imagination is a powerful tool here, because what we are capable of imagining can outstrip our discursive concepts. [13] Someone who has no technical knowledge of music, no concept of time signatures and pitches and so on, can nonetheless remember a tune by running through it in his head.[14] He can also identify the tune for someone else by a similarly applied method, that is, by singing the tune aloud. Here the tune is denoted "by displaying an instance of it for auditory sampling" (Foster 2000: 118). This is one illustration of the use of samples in communication.[15] Now, to utter a perceptual report is not to paint a (literal) picture. The speaker doesn't provide a sample to aid our

---

[12] This label is most associated with (Chisholm 1957) and (Jackson 1977). Martin (2010) emphasizes the role of implicit comparisons in our talk about the looks of things. The verb "look" has its own extensive literature, and it's commonly proposed that several meanings be distinguished even within its use to talk about visual appearances. See also Price (1932), Chisholm (1957), Jackson (1977), Dretske (1995), Martin (1997), and Travis (2004). Breckenridge (2007) maintains, against the majority, that "look" has a single meaning in sentences that describe visual appearances.

[13] If there are also, e.g., demonstrative concepts, then this claim is compatible with conceptualism about the contents of perception. See (McDowell 1994). The claim would then be true because of a limitation, from the perspective of communication, of demonstrative concepts: one cannot make an informative use of such a concept unless the hearer also is, or has been, acquainted with the feature it picks out.

[14] An example from Foster (2000), who calls this form of identification a "quotational designator": the performance denotes the tune "by displaying an instance of it for auditory sampling" (118).

[15] See also (Davidson 1968); Quine's (1969) remarks on "deferred ostension" (41); and Heal (1997) on "indexical predication" (153).

understanding of a perception or a hallucination. It's been suggested in discussions of empathy, however, that we sometimes use our own mental states, including imaginative states, as samples with which to think and learn about the mental states of others.[16]

Applying that idea in the present context explains our ability to understand sensory experiences not our own, and predict impacts on reasoning, emotions, and behaviour, consistent with the thesis that some aspects of the subjective are simply ineffable. Of course *we* are quite capable of understanding the facts in question, and also of speaking about them. The source of the expressive limit is that we cannot get such facts across in a way not reliant on the speaker's and hearer's capacity to take on the perspective being described, a capacity that exploits their own status as conscious and their own repositories of experience.

5.12 Some final observations

Criticizing the set of assumptions that generate Cartesian scepticism about the senses, Barry Stroud (2009) describes an alternative "conception of the mind and its contents, or an understanding of psychological facts in general, as being all of a piece with nonpsychological facts of many kinds." (568) So conceived, psychological facts are "not intelligible on their own, completely divorced from facts of the accompanying non-psychological world in which we understand them." (*ibid.*) This is the conception that we are driven towards if we take seriously the idea that there are some states of mind that we can only get a grip on by pretending; and it's the understanding that I claim is evinced by our ordinary ways of speaking about sensory experience.

The hypothesis that some of our talk about what people see, hear, feel, taste, and smell is an exercise in pretence has a number of distinctive payoffs. One is that it justifies our use of non-literal speech in a way that other accounts cannot. If what we mean by *Macbeth sees a dagger* is something to the effect of *it seems to Macbeth that he sees a dagger*, the main reason not to say the second thing straight out is that the first is marginally more concise. On my view, there are some facts whose acknowledgement *requires* an incursion into the non-literal: the subjective reality of hallucination is in one important sense ineffable.

Second, my view predicts the substitution-failure intuitions that perceptual reports can elicit, which have sometimes been taken for signs of intensionality.[17] In §5.6 I characterized the language of perception as economical, serving to describe the objective world and a perspective on it in one stroke. Part of what we typically glean from a perceptual report, then, is an idea of how things are for the perceiver. But different ways of identifying a single object are likely to bring to mind different images, different clusters of appearance-properties.[18] Since "Lois Lane sees Superman" sets up a certain expectation about the character of Lois's experience, it will tend to be uncooperative if, e.g., the person she was confronted with was wearing a suit and glasses. All the more so if the sole aim of the statement was to comment on her perspective, as will inevitably the case if we are ascribing a hallucination.

---

[16] (Walton 2015); Heal (1998) considers this possibility in connection with simulation theory.
[17] Notably by Anscombe (1965) and Brogaard (2018, chapter 6).
[18] See Braun and Saul (2002), especially pp. 25–6.

Third, by identifying certain uses of perceptual reports as a form of linguistic pretence, we can also account for the fact that the effect Moore first observed extends to such a wide range of syntactically- and semantically-disparate sentences. Notice, for example, that it's not just the verbs in perceptual reports that do work here: there is a striking reliance on talk about *things*—would-be spatio-temporal particulars—in our descriptions of sensory states, veridical and otherwise. An example to be understood in this light, where a perceptual verb doesn't even appear: "When I take my glasses off, everything is blurry." Of course it's not the case that things are blurry; a problem with my eyes has a consequence for my visual phenomenology. But how better to describe that consequence? Consider also in this light some of Wittgenstein's (1953) remarks on aspect-seeing. Describing a moment when the solution of a 'puzzle-picture' comes into view, when a hidden image coalesces, one can say: "Before, there were branches there; now there is a human shape." (196; my emphasis) That "the visual impression has changed" (*ibid.*) is brought out by speaking as though something in the world, the puzzle-picture itself, has undergone a transformation.

Finally, I note that the account I've defended here straightforwardly explains the possibility of expanding on a description of a hallucination in the following way:

> She sees a persimmon. It is on the table and is perfectly ripe.

The impression that we are speaking as though the putative object of perception exists is particularly strong with examples like this one. There is a close kinship between this series of claims—with its repeated anaphoric references connected to an indefinite that occurs in the scope of an attitude verb—and the phenomenon of intentional identity that we encountered in §4.2. The speaker here is *referring repeatedly to a single object*, i.e., making a series of identifications within a pretence about an object that exists to be seen only within that pretence.

This is also the moment to revisit the early debate about the perceptual verbs, assessing it with a new account of the two uses in hand. Moore and the sense-datum theorists (among whom we should also count Ayer for present purposes) are to be credited for their sensitivity to the relational character of perception and the relational structure of the reports we make in characterizing our sensory experiences. But these philosophers also make the mistake characteristic of the ontological strategy: accounting for the phenomenology of hallucination, and for the aptness of hallucination-ascribing reports, by conjuring an object and a relation of awareness *in reality*. I have situated them in a pretence instead.

Austin rightly insists on a single meaning for "see" and correctly identifies the second use of perceptual reports as non-literal. But he seems to view this position as less stable than it in fact is, as witnessed by his efforts to minimize the various examples in which the language of veridical perception is invoked to talk about hallucinations and illusions. There is no reason to fear examples of this kind. To the contrary, it's on the supposition that seeing is by definition a relation to ordinary things that descriptions of non-veridical experiences in the terms of veridical perception have their best justification. Located within a pretence of the kind that I have described here, the occurrence of such reports conforms to the thought that we conceive of perceptions as experiences whose character is supplied by the character of their objects.

Anscombe's approach is characteristic of the representational strategy. Her reaction to the appropriateness of identifying an object of sight even when there is no object in the "modern" sense is to attribute that appropriateness to a grammatical feature—to the fact that the verbs of perception are transitive—and to reclaim a sense of "object" that is free of existential implications. I would rather suggest that the verbs of perception are transitive because we understand perception to be a relation, and that these transitive verbs recur in reports of hallucination because we understand hallucination to be introspectively like perception. Readers familiar with Anscombe's paper may recall her evident aversion to the two figures she is responding to: the sense-datum theorist for whom our talk of "objects of sight" licenses the introduction of a new kind of entity, and the ordinary language philosopher who refuses to acknowledge that claims about how things appear using "see" are commonplace and legitimate. But the account I've pursued here allows us to side with the ordinary language philosopher on the question of verb meaning while siding with the sense-datum theorist on the legitimacy of both uses of perceptual reports—a compromise that is responsive to Anscombe's main objections, though it is not the one she pursues.

What about the talk of silvery specks and white dots that provoked the literature on perceptual identification? As I emphasized in discussing Austin's reaction to these cases, they involve neither hallucination nor illusion. "A silvery speck no bigger than a sixpence" is just an evocative description of the look that stars under what are, for us, normal viewing conditions (from a great distance, and through our planet's atmosphere). Hence the moral that Austin missed: these examples show how thoroughgoing our reliance on the procedure I have been drawing attention to really is. Even when there is something that we see (a star, a house, a decoy, twin A), we may prefer to say that we see another thing (a silvery speck, a white dot, a duck, twin B). We do this because the description we generate by invoking an object other than the object seen is as faithful or more faithful to the way that our perception strikes us.

As concerns the statement that *the silvery speck is a star*, and the relative infelicity of "that star is a silvery speck," we might compare an example of what Donnellan (1966, 281) calls the "referential" use of definite descriptions:

> The man over there drinking champagne is happy. (Kripke 1977, 256)

Donnellan observes that in evaluating this claim for truth or falsity, we are not especially concerned about whether the description is really met or by whom—it doesn't matter if the man that the speaker intends to talk about is actually drinking sparkling grape juice out of a champagne flute. What matters is that the speaker gets the hearer to fix on the intended object, and whether *that man* is happy. Why are we so cavalier about the definite description? Because, as Donnellan's label suggests, the point of the description was just to make a certain object salient. By that measure, it was well chosen.

Next we need to think about the kinds of contexts where statements like "That silvery speck is a star" or "that dot on the horizon in my house" appear. These are both claims that identify or categorize an object of joint visual attention: they say which thing or what kind of thing is being observed. Such comments are useful when it is *not already obvious* to the hearer that they are looking at a star or at their companion's house. In these circumstances, such *true* descriptions as "a star" and "my house" cannot be used to attract the hearer's attention to the object that one wants to talk about. "That silvery speck" and

"that dot on the horizon" exploit descriptions of appearance to draw attention to an object, and then a claim is made about the object so identified. Whereas the speaker in Kripke's example above is plausibly just making a mistake about the contents of the target's glass, the speaker in the "silvery speck" and "white dot" cases is deliberately prioritizing the looks of things over their true nature. Arguably, this is to be cashed out as a matter of pretending (speaking as though) an object is of a different nature than the speaker knows it to be.

This is just a sketch of a proposal, and I will not pursue its elaboration here. But note that if we interpret these sentences in the way that I propose to, they share a feature that Evans (1982, chap. 10) attributes to negative singular existentials. In denying the existence of Pegasus, Evans suggests, we make an *intra-pretence use* of what is in reality an empty name in order to make *an extra-pretence claim* about existence. With our "perceptual identification" sentences, too, there appears to be a shift within a single sentence from non-literal to literal speech, pretend-reference to genuine assertion.

5.13 Conclusion

We should expect the everyday understanding of perception to involve the kind of data that are available in everyday life, which in this case come in two main kinds: the first-personal data of experience, and the data from other perceivers who can corroborate or confute our own reports. Hence the notion of a way things can be for someone that is subjectively indistinguishable from seeing, but which lacks the latter's usual grounds; a state that we can expect to have the same kinds of impacts on behaviour and belief formation as perception does, but that lacks its epistemic value.

Hallucination-ascribing uses of perceptual reports have revealed themselves to be highly resistant to paraphrase. This is not a challenge to be overcome, but an indication as to the nature of the concept of hallucination, and of the way that we manage to talk about non-veridical sensory states. The concept is derivative, and our mode of description is indirect. The way that we go about filling out our characterizations of perceptual states of a given kind (seeing, hearing, etc.) is by describing the world around the subject. This is a strategy flexible enough to be mobilized in conditions where experience and reality come apart, because objects needn't be present, nor properties instantiated, in order for us to pretend that they are and imagine accordingly.

Chapter 6: Meeting the Challenge

It has been clear from the beginning that Prior's challenge is, at one level, easily met: just pick a claim to dispatch, and you have a consistent story to tell about the relationship between thinking of what exists and thinking of what doesn't. For a response to be at all satisfying, of course, there should also be good reasons to favour it over the other two options. But we might hope for still more. What's gripping about the series of claims that Prior presents is that *they all seem true*, even obvious. That's what gives the challenge the character of a trilemma. It would be nice if the impression of a forced choice between undesirable outcomes could also be dispelled, so that we might rest content with our response. As I've elaborated the disjunctive strategy, it can secure for us a resolution of that kind. In this final chapter, I first tie up some loose ends about my view and its implications. I then address a competing proposal about how Prior's challenge is to be put to rest. Finally, I revisit the three tempting claims with which we began.

6.1 Literalism

The availability of the position I have defended here is obscured from our view in two interconnected ways. One is through the influence of Cartesian epistemology, which I discussed in §2.1.2.1. In the intervening chapters, however, another reason for our unduly narrow sense of the possibilities has emerged: what we might call *literalism*.[1] By "literalism" I mean not only the tendency to assume that the literal meaning of an utterance is the one at issue, neglecting whatever falls under the heading of the pragmatic. I mean in addition a much more general conviction that everything that merits the title of "communication" is ultimately to be cashed out in terms of the mutual grasp of a proposition.[2] This philosophical idealization—amply justified in some contexts, for some purposes—is one that we are disposed to read back into reality.

In a Cartesian framework, as we have seen, one expects that *saying how the non-mental world is* and *attributing a mental state* should be two neatly-separable tasks; that psychological facts should admit of an independent characterization, so that no special problem is introduced by states of mind that make a radical departure from reality. This is part of what drives the pursuit of a true proposition about the subject's mind that could be the intended meaning—whether stated or conveyed—of those mental state ascriptions that I have proposed to understand as occurring internal to a pretence. But we are also driven in this search by what appear to be quite innocent assumptions about language. The convergence between these two lines of thought makes the conclusion I have sought to avoid feel all the more inevitable.

Here is a sketch of the literalist path to that conclusion. We start with an observation to the effect that uttering sentence *s* is conversationally appropriate, though it

---

[1] That there was an issue about the interpretation of language was clear from the start, hence my discussion of Anscombe's two uses of attitude ascriptions in chapter 2. But it's taken longer for a clear *target* to come into view.

[2] The point is not that there is a problem with defining a narrower concept of communication (as I have followed Evans in doing above). What I would object to is the thought that this is the only legitimate understanding of the notion, excluding the broader one.

has a salient interpretation on which it expresses a falsehood. We add that when declarative remarks are apt, it's in virtue of their being true. Stopping here, we conclude that *s* is ambiguous, or that our assumptions about its standing meaning need to be revised. Pressing onward, we may add that *if* the aforementioned principle appears to be violated, it's because we were wrong to fix on the literal meaning of *s*. Then we will posit a pragmatic mechanism and an alternative true claim that, we conclude, *s* conveys. Either way, we take it that *some truth is communicated*. When it comes to the mental state ascriptions that have sowed such confusion, that truth cannot be a claim about the relation between one or more subjects and a particular, since in the cases at issue (hallucination, illusion of singular thought) there is none. The most salient alternative is that we are communicating a truth which exclusively concerns the intrinsic state of the subject—a claim about a mental state she is in whether there is a particular thing that she thinks about or not.

The case against literalism to be extracted from this dissertation has several dimensions. We've taken note of the role that this assumption has played in causing debates about certain psychological ascriptions to stall out. There is also the simple fact that it causes us to neglect altogether the sort of account that I have explored here—an account that explains the communicative and epistemic payoff of certain speech acts other than by pointing to a content that these (really) express. Then there is the way that literalism entrenches, without actually lending support to, a view about the mind that we might wish to reject. But one can also cast doubt on this framing assumption in a more direct fashion, and I shall briefly do so now.

It is instructive to consider how the proposal that I draw from Walton and Evans contrasts, in its handling of some of the speech that we might qualify as "non-literal," with a familiar pragmatic paradigm. Here is Paul Grice's description, in "Logic and Conversation," of the process by which an implicature can be worked out:

> [The] speaker has said that *p*; there is no reason to suppose that he is not observing the maxims, or at least the Cooperative Principle; he could not be doing this unless he thought that *q*; he knows (and knows that I know that he knows) that I can see that the supposition that he thinks that *q* is required; he has done nothing to stop me thinking that *q*; he intends me to think, or is at least willing to allow me to think, that *q*; and so he has implicated that *q*. (1989; p.31)

As concerns the sorts of utterances I have been focused on throughout this dissertation, I am happy to follow this narrative up to a point. The speaker says that *p* (e.g., Macbeth sees—veridically perceives—a dagger). The context makes it clear that *p* is not true. The conversation is cooperative; the hearer knows that the speaker intends to inform and to speak truthfully. So the speaker cannot intend to communicate *p*. This provokes the search for an alternative. The account I defend deviates from Grice's at this stage, in that what vindicates the saying of *p* may not be the retrieval of an alternative proposition, *q*. No proposition is available; none suits the speakers communicative intentions and aims.

Against the literalist, I note that there are cases where this last model seems more appropriate that have nothing to do with giving expression to the subjective perspective. Consider the following example. The eponymous main character of the 2022 film *Tár* is a successful conductor, ruthless in both her work and her personal life. Having learned that her daughter, Petra, is being bullied at school, Lydia Tár tracks down the child

responsible. She then introduces herself as Petra's father. There is no clear candidate proposition conveyed by "I am Petra's father" that is also true in the context. To see why the remark is fitting, you are best off watching the movie up to that point. Even supposing, moreover, that the message could be adequately paraphrased, that paraphrase wouldn't have the same bite as the non-literal remark. What Lydia says is *more* apt than any literal truth would be. Her choice of communicative means is exactly right, though it's doubtful that the audience's understanding consists in the grasping of a proposition, left unarticulated, that she herself had in mind.

Donald Davidson (1978) takes a similar view of metaphor, and his remarks in that context will help to illustrate the claims and commitments associated with a rejection of literalism in a given sphere (a function they can serve irrespective of whether he is right to reject its application to metaphor). Like the attitude reports I am interested in, metaphor is a linguistic phenomenon that many would account for by appeal to the distinction between what is said and what is conveyed. One person tells another, "You are the cream in my coffee." The literal meaning is patently false, but the utterance is in order; we conclude that what's at issue must be some other content, true in the circumstances. What is that content? We might be offered a paraphrase like "you complete me," but this seems only to be a rough spelling-out of the message, and with some uses of metaphor we might not even get that far.

Some will persist in the attempt to track down a literal restatement of the metaphorical claim's content. Davidson proposes instead that we reconsider the approach that leads us to expect an exact paraphrase:

> There is... a tension in the usual view of metaphor. For on the one hand the usual view wants to hold that a metaphor does something no plain prose can possibly do and, on the other hand, it wants to explain what a metaphor does by appealing to a cognitive content—just the sort of thing plain prose is designed to express. [...] There is a simple way out of the impasse. *We must give up the idea that a metaphor carries a message, that it has a content or meaning* (except, of course, its literal meaning). (45; my emphasis)

> The central error about metaphor [...] is the thesis that associated with a metaphor is a cognitive content that its author wishes to convey and that the interpreter must grasp if he is to get the message. (46)

Davidson's alternative is to emphasize the way that the sentences in question are used, holding their literal meaning fixed—it's because of the use, not the meaning, that these count as cooperative contributions. (43) Metaphor, he suggests, ""runs on the same familiar linguistic tracks that the plainest sentences do" (43). This is the positive side of the rejection of literalism, that is, what such a rejection redirects us towards; it has manifested in the foregoing chapters as the view that the right pragmatic account of such claims as "She is worried about Santa Claus" and "Macbeth sees a dagger" is one that appeals to a distinct kind of speech act (a pretend assertion) rather than to a distinct content.

Though some attempts to paraphrase non-literal speech are *plainly* inadequate, this is not always so. There are candidates that I've rejected for which this is not the true, candidates that my reader may be genuinely tempted by: attitude ascriptions with explicitly general (existential) content, descriptions of experience in terms of "seems" or

"appears." It's because these strike many philosophers as satisfactory that I have had reason to advance various arguments and considerations in an attempt to unseat them. Davidson's comments draw attention to some additional ways of pushing back. He emphasizes that in rejecting the search for an exact paraphrase (a statement in literal terms of the content ostensibly conveyed by the utterance) one needn't deny the utility of various literal statements in the vicinity, for instance in helping the unsubtle hearer to understand what the speaker's point was (47). He also warns against the possibility that one is "[fastening] on the content of the thoughts a metaphor provokes and [reading] these contents into the metaphor itself" (45).

It's plausible that an error of this kind figures in the debate about perceptual identification that I discussed in chapter 5 (§5.2). Hinton discusses the relationship between

> "That speck of dust is an airplane." (1967a, 425)

and

> "What we see, looking a bit like a speck of something being wafted in the air . . . is the airplane." (*ibid.*)

Though a common way of speaking, the first of these is not obviously a literal truth. The second is. Hinton calls statements of the former kind "fragile" because we are "[reluctant] to lean on them" and have a "tendency under pressure to replace them" (*ibid.*). He calls statements of the latter kind "spotless," and offers the following characterization:

> When it is a question of repeating or not repeating, word for word, some such remark to someone who did not hear you the first time; or when you are reporting such a remark to a third party; or when you are explaining to a third party what you meant by 'I believe you're right,' said apropos of such a remark; or when you want to say that if the remark is correct then such and such a thing is the case; in situations like these, you tend to make such a replacement [of the fragile remark with a spotless one]. (*ibid.*)

The spotless statement, in these perceptual examples, will be one that "obviously neither entails nor presupposes the existence or subsistence in the relevant situational context of any speck, blob, blur, patch, disc, or whatnot in any sense of those words." (*ibid.*) In Hinton's terms, the error that Davidson diagnoses is a matter of taking the spotless replacement for a fragile statement to be the content that the fragile statement conveyed.

Hinton's description of the circumstances in which such replacements tend to occur brings out that the speaker is making a trade-off in moving from one kind of statement to the other. Grant for the sake of argument that the fragile statement is the most apt. Because it is not strictly true, its aptness comes at a cost: a chance that the hearer will take it literally and be confused or misled. The balance of risk and benefits is shifted in the circumstances that Hinton lists, in such a way as to favour plain expression. The pull of "seems" and "appears" statements, for example, stems in part from the fact that they tend to be true in those cases where we use perceptual reports to ascribe

hallucinations, as "I see a stick that looks bent" tends to be true when one uses "I see a bent stick" to report one's experience of a straight stick in water.[3]

So the perceptual report (fragile, on the use in question) is replaceable by another statement (spotless) for some purposes, e.g., to correct a misunderstanding:

> *A:* "And then, consumed by fever, Macbeth saw a dagger floating before him…"
> *B:* "A floating dagger? How can that be?"
> *A:* "Sorry, it merely *seemed to Macbeth* that he saw a dagger floating before him. He was hallucinating."

The error one is prone to falling into here is not recognizing that *A*'s reply *revises* the original statement for *B*'s sake.[4] One then takes the reply as offering the content that the unadorned perceptual report was supposed to convey, and which (if only it *had* been conveyed successfully) would have vindicated the speaker's choice of indirect means of communication. One takes the "seems" statement to be expressing what was meant all along. No: *A* has given up on her initial ambition, which was to have the hearer inhabit Macbeth's point of view. When she finds that this effort has sown confusion, she salvages the minimum necessary to keep the exchange on track; she characterizes Macbeth as the confused and delirious person that we know him, from the outside, to be.

Davidson's doubts about the accounts he opposes issue from a sense that metaphors are richer in significance, their interpretation more creative and more open-ended, than the accounts he opposes would seem to allow (1978, 31). My dissatisfaction with the literalist treatments of the attitude reports I have set out to explain is sparked in a couple of ways. One factor is the worry that these treatments appear adequate only against the background of an impoverished conception of thought and experience, the result, in turn, of a capitulation to radical scepticism (about which I say more in §6.3 below). Another factor is the fact that these uses of language are widespread, systematic, and also difficult to paraphrase in a way that replicates their precise effect. These features constitute a prima facie case for expecting that "Macbeth sees a dagger" or "She is afraid of Bigfoot" are fulfilling a special purpose, responsive to a particular need. The view I have defended here accords them just such a role.

6.2 An expressive limit

Connected to my rejection of literalism is the claim that some of the facts we acknowledge—subjective facts, specifically—are *ineffable*. That is, they are not describable in the same way that other, objective facts are describable, in literal terms. Ineffability gets invoked in different ways and for different reasons. There is the idea that the richness of our experience outstrips the descriptive resources of our vocabulary, or the claim that no amount of reading or study of the human visual system can convey *what it is like to see red* to the unacquainted. There is the Buddhist view that ultimate reality (as opposed to "conventional" or "historical" reality) is inexpressible, or Wittgenstein's talk of what is mystical and of "passing over in silence" (1922). While the ineffable by

---

[3] Though they won't *always* be true; see (Bourget 2017, 521–2) for some counterexamples.
[4] Compare the non-philosophical observer of chapter 5, §5.9

definition cannot be put into words, one can in various ways talk around it; in particular, one can pose and address such questions as:

- In what ways, if any, is this ineffable *x* nonetheless accessible to the minds of humans? In what ways, if any, is it describable or expressible?
- Why does this *x* resist (literal) description? What about the tool of language is unsuitable? What about the task makes it so?

The origin of the ineffability of some subjective facts, according to the view I have articulated, lies outside the mind of the subject: in the world's occasional failure to supply the state of affairs that we would need to use in order accurately to characterize the perspective in question.

That this should impose any sort of constraint is revealing as to the nature of the subjective itself. To be a subject is to occupy a point of view on the world. Hence the essentially perspectival character of facts about consciousness, which troubles their integration into our scientific worldview by one prominent measure of what that integration requires—namely, the faithful translation of perspectival claims into an objective, in the sense of "perspective-free," idiom. Hence, also, the power of pretence. The way to access perspectival facts is through the adoption of the relevant perspective—and that is bound to be an imaginative project most of the time.[5] The project is not, as it would be on a view of the mind as self-contained and therefore fundamentally inaccessible to outsiders, a wholly speculative and freewheeling one. The perspective is individual, but it is a perspective *on* the world that we all share. Recall that pretence as Walton defines gravitates around props, which are things in the world, property instances, real occurrences, psychological states, and so on. Pretending has us keyed into the world even amid efforts to engage with the idiosyncracies of a given point of view on it.[6] As we've seen, moreover, pretending can be existentially as well as "qualitatively" creative. These are just the features that one needs to reproduce in imagination the structure and the richness that the objective world provides to a course of experience.

6.3 Treading carefully towards a radical view

In chapter 2, I used Anscombe's description of what she calls the "intentional" and the "material" uses of attitude reports to slow down a line of thought—from the observation that we are given to utter the same string of words in two different contexts to the conclusion that there is a single kind of mental state that we are characterizing in both cases—that normally unfolds with undetectable swiftness. I don't share Anscombe's stopping point, and I reject some of her conclusions. But she brings it to our attention that

---

[5] The exceptions being cases where one can, for example, literally take over the spatial perspective of another person by moving to the place that they previously occupied, which is one source of insight into the sensory experiences of other people.

[6] That make-believe involves us with the world is a feature that Walton also highlights: "Sometimes, to be sure, imagining is a means of escape from reality, and we do frequently imagine what is not really the case. But even when we do, our experience is likely to involve the closest attention to features of our actual environment, not a general oblivion to it. Most imaginings are in one way or another dependent on or aimed at or anchored in the real world." (1990, 21).

before one asserts that thinking of *that tree* and thinking of Pegasus are the same sort of thing, one should start with the more cautious claim that we use the same form of language across two kinds of case. We refer to the subject, we select an attitude verb, and then we follow that verb with (in these examples) a singular term. In addition, our linguistic intuitions suggest that we are licensed in doing this whether or not the singular term that follows the verb actually refers to anything. *This* observation, and *not* the idea that at some level a single kind of mental state is involved in both cases, is the datum that any account of intentionality has to take as its starting point.

But this datum about language leaves us with more freedom than we have if we accept the "same sort of thing" claim straight off. It invites us to ask, "Why do we say the same thing across these two cases?" And in particular, "Why do we say the same thing while also realizing that the situations we are talking about are importantly distinct?" For no one would deny that we do recognize a difference, and have practical reasons to mark it. The lines we draw between perception and hallucination, between thoughts about Atlanta and Atlantis, between knowledge and belief, between what you think you're doing and what you're in fact doing—these all reflect our interest in keeping track of what is true of the world and what is merely projected on it.

Anscombe's own paper does not suggest a determinate account of her two uses (this is the stopping point that I want to press beyond). She tells us that there is a more fundamental, core use in which the noun phrase that follows the verb does refer, and also presents the object under a guise that the subject of the ascription herself thinks of it under. Then there is the derived use we make when the subject's cognition or experience has departed from reality in some way, where what we decide to say is entirely dictated by her perspective, i.e., by what that subject would affirm.

We can view the pretence-based account that I pick up from Walton and Evans as a way of filling in this structure. The key philosophical insight, already there in Anscombe's presentation, is that there is a class of statements about the mental that are parasitic on another class of statements; the issue was just how to elaborate on the nature of that parasitism. The hypothesis that we engage in pretence does that in a simple and elegant way. I emphasize in particular that it does not introduce a new set of meanings; it does not require us to posit any entities that are not obviously part of folk ontology (fictions, possibilia, the merely subsistent but non-existent, etc.); it doesn't impute to ordinary speakers an implausibly philosophical or theoretical conception of the relevant mental states (compare the exaggerated but illustrative simplicity of Fub's procedure in §4.2.4); it doesn't require the introduction of propositions or thoughts that are singular without being object-dependent (contra advocates of the representational strategy), nor of genuine thoughts that are merely putatively singular (contra a straightforwardly positive disjunctivism), working as it does instead with *singular propositions and thoughts internal to a pretence*.

When I initially presented Evans's Russellianism about singular terms and singular thoughts in chapter 3, I introduced it as it tends to strike us on the first read, namely, as flying in the face of some basic features of the empty case. The approach I have taken throughout this dissertation, however, has us arrive at this radical picture from a relatively uncontentious starting point: some observations about what we say under what conditions, what we are trying to do in saying those things, and what success at that

task would look like. It becomes clear that the internalist end point is the product of a number of over-hasty moves that can and should be reconsidered. Among them:

- assuming that our talk about the mental, if it seems correct, is correct because literally true,
- assuming that if it's not literally true, it at least conveys a proposition that is true,
- assuming that subjective indiscriminability entails sameness of kind,
- assuming that the only reason we would have for saying "sees" in cases of perception and hallucination, or "loves" in cases involving real things and made-up ones, is that we think seeing [something] and loving [something] are states that can obtain regardless of the existence of what they purport to be about.

Understanding that we were guided towards internalism in this way is important for assessing the outcome of my arguments up to this point. One might suppose that we are left with two competing hypotheses (common-kind, disjunctive), both consistent with the evidence, such that choosing between them is a further task. That's not the case. We had a shared starting point that we seemed to be forced away from by philosophical reflection. But if the moves that led us to abandon the common-sense picture were based on a mistake, then showing that the evidence admits of an interpretation on which it is perfectly consistent with the common-sense picture doesn't leave us in a neutral position relative to the two kinds of account. If we notice that we made a mistake in our assessment of Prior's *(c)*, and we retrace our steps, we are simply led back to *(a)*: thinking of Y is a relation to Y when Y exists.

6.4 A competing solution

Prior's challenge begins with the common-sense idea that some of our thoughts are relations to the particular things that they concern. As one progresses through the subsequent claims, one is invited to conclude that this piece of common sense is, if not entirely mistaken, at least in need of a revision. This is the style of response to Prior that R.M. Sainsbury (2018) recommends as he makes the case for what I've called the "representational strategy." By noticing how this proposal falls short, we can better appreciate the distinctively satisfying character of a response in terms of the disjunctive strategy.

When one confronts Prior's three jointly inconsistent claims, one might wonder whether some sort of shift is occurring as one moves between them—a shift that could explain why they each individually seem true. Sainsbury offers an explanation of this kind: the appearance of an intractable problem issues from a failure to keep apart two ways that an intentional mental state might qualify as relational. One way of being a relation is common, Sainsbury suggests, to all intentional states: these are all fundamentally *relations to representations*. The other kind of relationality sorts the empty mental states from the non-empty ones: it's a matter of whether the subject's representation is itself related to anything. Thinking about Pegasus and thinking about Obama are both relations in the first way; only thinking about Obama is a relation in the second way.

Why is this view not vulnerable to the kind of objection that we saw Anscombe raise (in chapter 1, §1.5.1) that intentional objects are *ideas*? A first problem with that claim was the introspectively evident difference between thinking of Pegasus and thinking of an idea of Pegasus; the second, closely related, was the fact that we already acknowledge such a thing as thinking of an idea of Pegasus, which we treat as distinct from what we call "thinking of Pegasus" (period). But Sainsbury, in line with the general spirit of contemporary representational approaches, proposes that the relation to a representation is both phenomenologically and linguistically "covert" (148). That intentional mental states have this nature is an empirical discovery on the part of cognitive scientists; it's not given to the subject, discernible *a priori*, or capable of being read off from the way that we talk about the attitudes in question.

This response fails, I think, for the very reason that it manages to sidestep Anscombe's objection: the fact that one of the relations Sainsbury appeals to is covert makes it unsuitable for the explanatory role he assigns it. The intended payoff of Sainsbury's move is that we can have it both ways: all intentional states are relational in one sense, while in another sense only some are. But the challenge in Prior's challenge is, as above, not to come up with an account of intentionality that avoids contradiction; he himself points us towards three good ways of doing just that. What Prior distinctively emphasizes, besides the fact that one of *(a)*, *(b)*, or *(c)* must be abandoned, is the undesirability of this resolution. All three claims are intuitively appealing, and persistently so. This tells us something about what is required to put the challenge to rest. What's required, beyond mere consistency, is an explanation of *why* we are drawn to all three claims even though one of them must be false.

Sainsbury purports to supply such an explanation. We find ourselves in a bind because we fail to distinguish two forms of relationality, which categorize intentional states in different ways. But the tension on display in Prior's challenge—is *thinking of Y* a relation or is it not—arises from the most superficial reflection on what it is to think abouts a thing. That tension should, then, itself have its source in something relatively superficial: something about the everyday conception of these states, their phenomenology, the way we habitually talk about them. A covert relationality common to all intentional states does not fit the bill. Since Sainsbury's account of what draws us to *(a)* is not plausible, he fails to explain what is wrong about our embrace of this claim. He does not manage to diagnose *(a)* in such a way as to account for its appeal while also recommending its rejection.

Yet there is a different kind of shift that occurs when we contemplate the three claims that *can* be called upon to make sense of our initial indecision, the feeling of being in a bind. Rather than a change in meaning characteristic of equivocation, it's a change in the perspective that we take up: the perspective of the subject, or a perspective on the subject. Our failure to hold these apart, and to recognize exactly what is required to respect the subjective and the objective facts, is what provokes the feeling that any move we make will come at great cost.

From the third-person point of view, certain contrasts among mental states are highly salient: those between perception and hallucination, between knowledge and belief, between successful and failed demonstrative identification, for example. These contrasts do not always manifest themselves within the subject's point of view. In cases where they do not, the way to take up the subject's perspective is to ignore them. Because

of the asymmetrical relationship between at least some of these states, however—*perceiving* and *thinking of something* being the positive phenomena that we have an independent grip on, while *hallucination* and *merely apparent singular thought* are characterized by reference to these—ignoring the contrasts does not mean retreating to a more basic common element that was responsible for the subjective reality all along. It means imaginatively reconstructing the world from the perspective of the subject in question, a project that directs our attention outwards rather than inwards.

The insight that we all have, as subjects ourselves, into the subjective perspective; the ease with which many of us adopt perspectives other than our own; these are the sources of our desire to affirm *(c)*. Asked to consider whether thinking of Y is a relation to Y, we are pushed into a different frame of mind: the third-personal point of view that encompasses the subject, her intentional relations, and the objects of these relations. From this standpoint it is obvious that what we appealed to in the good case (a relation between the subject and the object of her thought) is not a factor in the bad case (when one relatum is missing). Continuously engaging one perspective and then another according to the claim at issue, we are taken in by each of the claims in turn.

6.5 The status of (a), (b), and (c)

Prior supplied us with a rough, working characterization of some claims we are committed to. By this point we've found it necessary to question and adjust them in various ways. I close by laying out, explicitly and in one place, what I propose is the reality behind the three claims.

6.5.1 "X's thinking of Y constitutes a relation between X and Y when Y exists"

On the implementation of the disjunctive strategy that I have elaborated and defended, a singular thought about a particular—a thought whose object is determined relationally rather than descriptively—is *fundamentally* a relation between the subject and that particular. I am now recalling my childhood home; the correct account of this mental state is the one that bridges the Cartesian as well as the internalist boundaries between the mental and the non-mental, the former delimiting these by the range of introspection and the latter by the physical boundaries of my brain or body. Contra the representational strategy, this "broad" condition is not a composite of a mental state (a relation to a representation, say) and a further fact about how I am embedded in my environment (a relation between the representation and my childhood home). The *subject-object relation* is primitive.

It has been a recurring theme that while the representational and the disjunctive strategies both give a role to particulars in their accounts of the good case, they do so in different ways. Here is one more variation on that theme. The question of what distinguishes a description of our experience and cognition in purely qualitative terms from a description that makes *essential* reference to particulars is one of Strawson's (1959) preoccupations in *Individuals*, and he addresses it by reference to the phenomenon he calls "reidentification": the possibility of coming across quantitatively the same thing again (31-2). Consider your experience as you look at a tree, blink, and look at it a second time. We might describe the experience you have when your eyes are open in the following two ways:

You see *that tree* (demonstrating).

You see something green, irregularly shaped, etc.

The first of these descriptions makes reference to a particular while the second does not. Yet it may seem that there is no important difference between them. As applied to the same perceptual episode, they report the same *objective* situation. The "something green, irregularly shaped, etc." of the second description will be identical with the tree we make demonstrative reference to in the first description. From the objective point of view, that is, both are descriptions of an encounter with a particular thing.

Strawson's point about re-identification draws attention to the fact that these are nonetheless descriptions of different *subjective* situations, of different courses of experience. The purely qualitative description doesn't provide for the re-identification of particulars. Something is required of the world to make it the case that I see *the same tree again*, or think of *the same home I was missing last week,* which is not required to make it the case that I see something green and irregularly shaped, or think of a home of such-and-such description, two times in succession.

The description of our psychological lives that makes reference to particulars harmonizes the subjective point of view with the objective one. The tendency among philosophers, which the representational strategy exemplifies, is to drive the two apart. To characterize the subjective point of view, one supposes, is to move away from the world and into "the mind itself." It's a matter of describing something that remains unchanged across a range of scenarios, scenarios that one is not in a position to distinguish on an introspective basis: perceiving and hallucinating, knowing and falsely believing, acting or merely seeming to, being an embodied creature or a brain in a vat, being shut up in a stove or deceived by an evil demon.

These skeptical reflections lead us away from our original assessment of our own situation and towards an internalist conception on which one needn't look beyond a subject's brain or body to determine what mental state she is in. They lead us to give up on the idea that particulars themselves can figure in our experiences and our thoughts. My defence of the disjunctive strategy constitutes a defence of the view that we do not need to give this up.

6.5.2 "X's thinking of Y is not a relation between X and Y when Y does not exist"

In chapter 2 (§2.1.1), we found Prior's *(b)* in need of a minor revision: to avoid prejudging any questions about the nature of the empty case, we might say something to the effect that the mental state we call "thinking of Y," where "Y" does not refer, is not a relation to Y. In this way we remain open to the possibility that, as Evans and McDowell claim, a merely apparent singular thought is in fact no thought at all. This second of the three claims is one that the disjunctive strategies and the representational strategies converge on, setting them against the ontological strategy; I shall now make some final comments on that last approach.

I first want to highlight what I think are mistaken reasons for rejecting the ontological strategy and the various accounts throughout the history of philosophy that, in the same fashion, endeavour to explain the mind by positing a new kind of thing—

substance dualism and sense-datum theory being two notable examples. The grounds I would resist are charges of ontological profligacy, or alarm at the strange and unusual nature of the entities or the substance in question. I do not deny that these reactions take account of real costs associated with the ontological strategy and its associates. The problem is that rejecting these views on the basis of their most obvious weaknesses, we move on without having conducted a full audit of them. As a consequence, we may leave in place some more pernicious errors.

This is particularly worrisome when a view of the kind in question does not simply exist alongside others, but represents a past stage of our philosophical trajectory. Such is the relationship between substance dualism and internalist physicalism, or between sense-datum theory and intentionalism: the latter are the results of our moving on from positions that became unpalatable *on ontological grounds* once it became generally accepted that the physical world is causally closed.[7] As has emerged throughout the preceding chapters, these views inherit what are arguably the deepest problems with their more metaphysically extravagant predecessors.

My basic case against the ontological strategy takes a different form: it's that an elaboration of the disjunctive strategy that appeals to pretence in its account of the empty case better achieves what the advocate of the ontological strategy, on my most charitable interpretation of their motives, is trying to do. They share my conviction that relations to things need to figure in our account both of the cases where (as we would first have put it) the object of thought exists, and of the cases where it does not exist. This best makes sense of the subject's point of view, the apparent possibility of coordination among several subjects (brought to our attention by Geach's Hob-Nob example), the way that we describe the mental states in question, and so on.

These observations are incorporated, in the context of the ontological strategy, via the claim that that strictly there is no such thing as what we have called the "empty" case— only a distinction among kinds of object of thought or among the grades of reality that they possess. I agree that there is no getting away from relations to things in our understanding and our characterization of even those mental states I have deemed defective; I agree that there is a sense in which even the empty case is object-dependent. But this is not the straightforward dependence of thought on its object. It's rather that merely putative singular thought is dependent on (derivative of, characterized by reference to) genuine singular thought, which *in turn* is dependent on the object thought about. These remarks can be parlayed into an error theory of sorts: the advocate of the ontological strategy mistakes pretending that there is an object for there being an object, pretending to refer to an object for actually referring to one.

6.5.3 "X's thinking of Y is the same sort of thing whether Y exists or not"

---

[7] The reason this made the views in question obsolete is that the mind (as conceived by Descartes) and sense-data (as conceived by the early sense-datum theorists) were both supposed to be at once non-physical *and* implicated in causing physical events, e.g., a person's reaching out for a glass of water. For insight into the perspective of philosophers whose career preceded this shift in our understanding of reality, one might peruse such papers as H.H. Price's "Some Philosophical Questions about Telepathy and Clairvoyance" (1940).

The disjunctive response to Prior's challenge requires us to deny a more determinate restatement of this third and undeniably appealing claim, namely that X's thinking of Y is a mental state of the same fundamental kind whether Y exists or not. This would be a genuine cost, a true case of "giving something up," if a belief in that more determinate claim was the source of the intuitive appeal of *(c)*. That is not the case. I've identified two other sources, which were the focus of chapter 2: there are the things that we say, and there is the understanding we have of the subjective perspective.

In each case there was a first phase of argument which aimed merely to establish the *compatibility* of these observations with the commitments of the disjunctive strategy. For the observations about language, I drew inspiration from Walton's point about the role that pretence might play in his imagined movie-goer's claim to have been "terrified" of the slime. One reason for Charles to say that he is afraid is that he is afraid. Another reason for Charles to say that he is afraid is that he is pretending to be afraid. Advocates of the ontological and representational strategies would claim as evidence for their own views the fact that we go on using locutions like "thinking of" and "wanting" whether really is such a thing, or such a *kind* of thing, as the subject takes herself to be related to in thought: as evidence that we count so-called "empty" cases as cases of singular thought. I propose in response that in the one case (singular thought on the basis of acquaintance) we are speaking literally, and that in the other case (merely apparent singular thought in the absence of acquaintance) we are either mistaken or pretending. This gets the disjunctive strategy on a par with the other responses to Prior's challenge, as far as consistency with the core linguistic evidence is concerned.

For the observations about the subjective perspective, pretence served in a different way. Here, my opponent would object as follows. "Suppose I take up the perspective of the child who believes in Santa Claus; I consider the relationships between her mental states, predict how she will act, and so on. I don't imagine the child's subjective life to be radically different from the way that mine is when I think about a real person. I predict that she will act just as someone with one of these ostensibly object-dependent singular thoughts will act." In fact my view predicts exactly this; it simply predicts it on a different basis than the advocate of the representational or ontological strategy does. As I noted in chapter 3, the child is objectively not having a singular thought, and yet "subjectively" she is: the right account of her point of view is an account that makes reference to an individual as the object of her singular attitude. There being no such individual and no such attitude, the way to take up her perspective is to engage in an existentially creative pretence for which informational states serve as props.

The second phase of argument aimed to establish that the observations that drive us to affirm *(c)* are not merely compatible with but in fact *better explained* by the disjunctive strategy than by the alternatives. In the case of the observation about ordinary mental state ascriptions, this was revealed by the shortcomings of the standard, literalist interpretations of the various statements we considered in chapters 3, 4 (specifically §4.2), and 5, which sometimes failed to predict particular features of the use in question (e.g., the fact that perceptual reports only have a specific reading, even as invoked to characterize hallucinations). Why do we insist on saying things like, "Macbeth sees a dagger"? The view I defend provides an explanation of that behaviour that is more robust and specific than the one that we must turn to if we take such utterances to be legitimated by the communication of a true proposition. Instead of appeals to such general

motivations as a desire to avoid prolixity or to add some colour by venturing beyond the strictly literal, we have an account of these utterances as responsive to the character of the subjective realities that elicit them in a way that dominant philosophical accounts of those same realities have failed to be. They are responsive, specifically, to the way that particulars figure in our subjective lives.

I want to close out this chapter and the dissertation by setting out as clearly as possible the analogous move for the subject's point of view, a matter of explaining how the disjunctive strategy's account of what it is like for a subject in the empty case can be shown to better explain the judgment of subjective sameness than its internalist rivals do. Recall that views of the latter kind reject Prior's *(a)* in order to hold onto a version of *(c)*, i.e., that the subject is in the same fundamental kind of mental state whether the object of her thought or experience exists or not. One of the initial motivations for maintaining *(c)* is the innocent or naïve subject's verdict that the cases are introspectively the same. Together with the idea that such judgments are revealing as to the nature of the mental states in question, this verdict is metaphysically significant, and *(c)* cannot be given up.

In short, the internalist is motivated to give up *(a)* by a desire to be faithful to the subjective perspective. She commits to theorizing that perspective *independently*, with the relations to objects to be added back in later. So, she will have an account of sensory experience, and some of those states will get to be perceptions in virtue of the right relation obtaining between the subject's sense organs and the environment. She will have an account of belief, and some of those states will get to be knowledge because of their pedigree. She will have an account of singular thought as distinguished by the kind of mental representation in question (a mental file, an individual concept, etc.), and some of those thoughts will really be about particular things, because of how the representations are connected to those things. Her account of the good case, and typically also of the associated concepts, will be conjunctive.

The unintended and frequently unrecognized result is a falsification of our subjective life as we knew it to be at the outset—of what we first set as the target of our theories and our descriptions.[8] What we began with was the observation that some mental states are subjectively, psychologically *just like* thoughts about particulars; and we were taking the latter to be relational until we were prompted by some philosophical reflections to say otherwise. The disjunctive strategy's characteristic claim, that thinking of Y is *fundamentally* a relation, admittedly makes it look like something is being added that was not already there in the everyday conception of these states, as though a common-sense claim has been replaced by a metaphysical thesis. But the qualification is just a defensive move against the advocate of the representational strategy, whose strategizing about how the relation might be broken down made it seem like there was something incomplete or vague about the original claim.[9]

---

[8] There are ways that this falsification might be thought to come out in the wash, which I don't explore here; e.g., in attempts to make sense of the way that an agent's psychology evolves over time (Williamson 1998, sec. 3) or in predictions about the likelihood of certain courses of action (Williamson 2002, 62).

[9] It is rather like the situation produced by the sense-datum theorists, who, by suggesting that we perceive sense-data directly and ordinary things indirectly (a way of speaking about perception with no precedent in ordinary language, as Austin complains (1962, 15), forced their opponents to

In the attempt to make sense of how the empty case could fail to be distinguishable by introspection from a case of acquaintance-based thought about an individual, the representational strategy has us accept that one's mental state in the *good case* is not a relation to the object of thought (instead, the mental state is a relation to a representation, the relation between the representation and the object being extra-mental). According to the representational strategy, then, the empty case is *not subjectively like* having a thing (literally) in mind; it is subjectively like being in an object-independent representational state. In other words, advocates of the representational strategy mischaracterize the good case in an effort to do right by the empty case; but their endorsement of *(c)* ensures that the mischaracterization spreads to the empty case, thus undermining their own aims. The outcome is that contrary to what is widely believed, it is not the disjunctive picture, but the common-factor one, that sets itself against the most obvious truths about thought and experience.[10] The disjunctive strategy preserves the integrity of the good case by according the empty case a derivative status, parasitic on the state of affairs that we routinely mistake it for; in the end, this is what it takes to preserve the integrity of the empty case as well.

A common misconception about views like the one I have defended here is that they are motivated by a sheer refusal to part with a naïve picture of our place in the world, where that parting is avoided by means of the philosophical equivalent of a legal loophole (e.g., that subjective indistinguishability doesn't *entail* sameness of kind). This is why I have been concerned to show that an account of singular thought as object-dependent is not just an *available* position, but that it is superior to the alternatives by a wide variety of measures. The moral of the argument just outlined, in particular, is that the disjunctive strategy should be attractive to anyone with the kinds of concerns about fidelity to the subjective perspective that first manifested themselves as objections to that very approach.

---

describe their own view as one on which we "directly perceive" ordinary things, and perhaps even to self-identify as "direct realists."

[10] Compare McDowell on the Cartesian's description of our experience: "Ironically, when reverence for the authority of phenomenology is carried to the length of making the fact that internal configurations are indistinguishable from the subject's point of view suffice to establish that those configurations are through and through the same, the upshot is to put at risk the most conspicuous phenomenological fact there is." (1987, 242–243)

References

Anscombe, G. E. M. 1965. “The Intentionality of Sensation: A Grammatical Feature.” In *Analytic Philosophy*, edited by R. J. Butler. Blackwell.

———. 1975. “The First Person.” In *Mind and Language: Wolfson College Lectures 1974*, edited by S. Guttenplan. Oxford University Press. Reprinted in *Metaphysics and the Philosophy of Mind: The Collected Philosophical Papers of G. E. M. Anscombe*, 21–36. Basil Blackwell.

Armstrong, D. M. 1968. *A Materialist Theory of the Mind*. Routledge.

Asher, N. 1986. “Belief in Discourse Representation Theory.” *Journal of Philosophical Logic* 15(2): 127–89.

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

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

———. 1956. *The Problem of Knowledge*. Penguin.

———. 1964. “Names and Descriptions.” In *The Concept of a Person*, 129–61. Macmillan.

Bar-On, D. 2004. *Speaking My Mind: Expression and Self-Knowledge*. Clarendon Press.

Barwise, J. 1981. “Scenes and Other Situations.” *Journal of Philosophy* 78(7): 369–97.

Barwise, J., and J. Perry. 1983. *Situations and Attitudes*. MIT Press.

Berto, F. 2008. “Modal Meinongianism for Fictional Objects.” *Metaphysica* 9: 205–18.

———. 2011. “Modal Meinongianism and Fiction: The Best of Three Worlds.” *Philosophical Studies* 152(3): 313–34.

Bourget, D. 2017. “Intensional Perceptual Ascriptions.” *Erkenntnis* 82: 513–30.

———. 2019. “Implications of Intensional Perceptual Ascriptions for Relationalism, Disjunctivism, and Representationalism about Perceptual Experience.” *Erkenntnis* 84(2): 381–408.

Braun, D. 1998. “Understanding Belief Reports.” *Philosophical Review* 107(4): 555–95.

Braun, D., and J. M. Saul. 2002. “Simple Sentences, Substitutions, and Mistaken Evaluations.” *Philosophical Studies* 111(1): 1–41.

Breckenridge, W. 2007. “The Meaning of ‘Look.’” DPhil dissertation, University of Oxford.

Brentano, F. (1874) 1995. *Psychology from an Empirical Standpoint*. Edited by L. McAlister. Introduction by P. Simons. Routledge.

Brewer, B. 2011. *Perception and its Objects*. Oxford University Press.

Broad, C. D. 1952. “Some Elementary Reflexions on Sense-Perception.” *Philosophy* 27: 3–17.

Brogaard, B. 2015. “Perceptual Reports.” In *Oxford Handbook of the Philosophy of Perception*, edited by M. Matthen, 237–254. Oxford University Press.

———. 2018. *Seeing and Saying: The Language of Perception and the Representational View of Experience*. Oxford University Press.

Burge, T. 1974. “Truth and Singular Terms.” *Noûs* 8(4): 309–25.

Burnyeat, M. 1982. “Idealism and Greek Philosophy: What Descartes Saw and Berkeley Missed.” *Philosophical Review* 91(1): 3–40.

Butterfill, S. A., and I. A. Apperly. 2013. “How to Construct a Minimal Theory of Mind.” *Mind and Language* 28(5): 606–37.

Byrne, A. 2001. “Intentionalism Defended.” *Philosophical Review* 110: 199–240.

Byrne, A., and R. Manzotti. 2022. “Hallucination and Its Objects.” *Philosophical Review* 131 (3): 327–359.

Camp, J. L. 2002. *Confusion: A Study in the Theory of Knowledge*. Harvard University Press.

Campbell, J. 1987. “Is Sense Transparent?” *Proceedings of the Aristotelian Society*, n.s., 88: 273–92.

———. 2002. *Reference and Consciousness*. Oxford University Press.

Carroll, N. 1990. *The Philosophy of Horror: Or, Paradoxes of the Heart*. Routledge.

Carruthers, P. 1987. “Russellian Thoughts.” *Mind* 96(381): 18–35.

———. 2005. *Consciousness: Essays from a Higher-Order Perspective*. Oxford University Press.

Cartwright, R. J. 1987. *Philosophical Essays*. MIT Press.

Chisholm, R. M. 1957. *Perceiving: A Philosophical Study*. Cornell University Press.

Chomsky, N. 1995. “Language and Nature.” *Mind* 104(413): 1–61.

Clarke, T. 1965. “Seeing Surfaces and Physical Objects.” In *Philosophy in America*, edited by M. Black, 98–114. Allen & Unwin.

Crane, T. 2006. “Brentano’s Concept of Intentional Inexistence.” In *The Austrian Contribution to Analytic Philosophy*, edited by M. Textor, 28–43. Routledge.

Crimmins, M. 1998. “Hesperus and Phosphorus: Sense, Pretense, and Reference.” *Philosophical Review* 107(1): 1–47.

Currie, G. 1990. *The Nature of Fiction*. Cambridge University Press.

D'Ambrosio, J. 2022. "An Empirical Solution to the Puzzle of Macbeth's Dagger." *Erkenntnis* 87(3): 1377–414.

Davidson, D. 1967. "The Logical Form of Action Sentences." In *The Logic of Decision and Action*, edited by N. Rescher. University of Pittsburgh Press.

———. 1968. "On Saying That." *Synthese* 19(1–2): 130–46.

———. 1978. "What Metaphors Mean." *Critical Inquiry* 5(1): 31–47.

Dennett, D. C. 1968. "Intentional Systems." *Journal of Philosophy* 65(4): 87–106.

Donnellan, K. S. 1966. "Reference and Definite Descriptions." *Philosophical Review* 75(3): 281–304.

Dretske, F. 1969. *Seeing and Knowing*. Routledge.

———. 1981. *Knowledge and the Flow of Information*. MIT Press.

———. 1995. *Naturalizing the Mind*. MIT Press.

Dummett, M. 1973. *Frege: Philosophy of Language*. Harvard University Press.

———. 1996. "Existence." In *The Seas of Language*, 277–307. Oxford University Press.

Evans, G. 1982. *The Varieties of Reference*. Edited by J. McDowell. Oxford University Press.

Findlay, J. N. 1961. *Values and Intentions*. George Allen & Unwin.

Fish, W. 2009. *Perception, Hallucination, and Illusion*. Oxford University Press.

Fodor, J. A. 1975. *The Language of Thought*. Thomas Y. Crowell.

Forbes, G. 2000. "Objectual Attitudes." *Linguistics and Philosophy* 23(2): 141–83.

———. 2006. *Attitude Problems*. Oxford University Press.

Ford, A. 2011. "Action and Generality." In *Essays on Anscombe's 'Intention'*, edited by A. Ford, J. Hornsby, and F. Stoutland, 76–104. Harvard University Press.

Foster, J. 2000. *The Nature of Perception*. Clarendon Press.

Frege, G. 1956. "The Thought: A Logical Inquiry." Translated by A. M. Quinton and M. Quinton. *Mind* 65(259): 289–311.

———. 1979. *Posthumous Writings*. Edited by H. Hermes, F. Kambartel, and F. Kaulbach. Translated by P. Long and R. White. Blackwell.

French, C. 2013. "Perceptual Experience and Seeing That p." *Synthèse* 190(10): 1735–751.

Friend, S. 2014. "Notions of Nothing." In *Empty Representations*, edited by M. García-Carpintero and G. Martí, 307–332. Oxford University Press.

———. 2017. “Fiction and Emotion.” In *The Routledge Handbook of Philosophy of Imagination*, edited by A. Kind, 217–29. Routledge.

Geach, P. 1967. “Intentional Identity.” *Journal of Philosophy* 64(20): 627–32.

———. 1980. *Reference and Generality*. 3rd ed. Cornell University Press.

Goodman, N. 1949. “On Likeness of Meaning.” *Analysis* 1: 1–7.

Goodman, R., J. Genone, and N. Kroll, eds. 2020. *Singular Thought and Mental Files*. Oxford University Press.

Greenspan, P. 1980. “A Case of Mixed Feelings: Ambivalence and the Logic of Emotion.” In *Explaining Emotions*, edited by A. O. Rorty, 223–50. University of California Press.

Grice, H. P. 1969. “Vacuous Names.” In *Words and Objections: Essays on the Work of W. V. Quine*, edited by Donald Davidson and Jaakko Hintikka, 118–45. D. Reidel.

———. 1989. *Studies in the Way of Words*. Harvard University Press.

Grzankowski, A. 2016. “Attitudes Towards Objects.” *Noûs* 50(2): 314–28.

Hampshire, S. 1964. “Identification and Existence.” In *Contemporary British Philosophy: Personal Statements, Third Series*, edited by H. D. Lewis, 189–208. George Allen & Unwin.

Hampshire, S, and P. F. Strawson. 1961. “Perception and Identification.” *Aristotelian Society Supplementary Volume* 35: 81–120.

Harman, G. 1990. “The Intrinsic Quality of Experience.” *Philosophical Perspectives* 4: 31–52.

Heal, J. 1997. “Indexical Predicates and Their Uses.” *Mind* 106(424): 619–40.

———. 1998. “Understanding Other Minds from the Inside.” *Royal Institute of Philosophy Supplements* 43: 83–99.

Heim, I. 1982. “The Semantics of Definite and Indefinite Noun Phrases.” Doctoral dissertation, University of Massachusetts.

Higginbotham, J. 1983. “The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics.” *Journal of Philosophy* 80(2): 100–27.

Hintikka, J. 1969. “On the Logic of Perception.” In *Models for Modalities*, edited by D. Davidson et al., 151–83. Reidel.

Hinton, J. M. 1967a. “Perception and Identification.” *Philosophical Review* 76(4): 421–35.

———. 1967b. “Visual Experiences.” *Mind* 76: 217–27.

Hoek, D. 2018. “Conversational Exculpature.” *Philosophical Review* 127(2): 151–96.

Hume, D. (1739) 1965. *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge. Reprint, Oxford University Press.

Ingarden, R. (1931) 1973. *The Literary Work of Art*. Translated by G. Grabowicz. Northwestern University Press.

Ishiguro, H. 1967. "Imagination." *Proceedings of the Aristotelian Society, Supplementary Volumes* 41: 19–56.

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

Jeshion, R. 2010. "Singular Thought: Acquaintance, Semantic Instrumentalism, and Cognitivism." In *New Essays on Singular Thought*, edited by R. Jeshion, 105–40. Oxford University Press.

Kamp, H. 1981. "A Theory of Truth and Semantic Representation." In *Formal Methods in the Study of Language*, edited by G. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof, 277–322. Mathematisch Centrum Amsterdam.

Kaplan, D. 1989. "Afterthoughts." In *Themes from Kaplan*, edited by J. Almog et al., 565–614. Oxford University Press.

King, J. C. 1993. "Intentional Identity Generalized." *Journal of Philosophical Logic* 22(1): 61–93.

———. 2020. "Singular Thought, Russellianism, and Mental Files." In *Singular Thought and Mental Files*, edited by R. Goodman et al., 73-106. Oxford University Press.

Kripke, S. 1972. *Naming and Necessity*. Basil Blackwell.

———. 1977. "Speaker's Reference and Semantic Reference." *Midwest Studies in Philosophy* 2(1): 255–76.

———. 2013. *Reference and Existence: The John Locke Lectures*. Oxford University Press.

Kroon, F. 1994. "Make-Believe and Fictional Reference." *Journal of Aesthetics and Art Criticism* 52(2): 207–14.

Lamarque, P. 1981. "How Can We Fear and Pity Fictions?" *British Journal of Aesthetics* 21: 291–304.

Levy, Y. 2013. "Intentional Action First." *Australasian Journal of Philosophy* 91(4): 705–18.

Lewis, D. 1968. "Counterpart Theory and Quantified Modal Logic." *Journal of Philosophy* 65: 113–26.

———. 1978. "Truth in Fiction." *American Philosophical Quarterly* 15: 37–46.

———. 1983. "Individuation by Acquaintance and by Stipulation." *Philosophical Review* 92(1): 3–32.

———. 1986. *On the Plurality of Worlds*. Blackwell.

Lockwood, M. 1971. “Identity and Reference.” In *Identity and Individuation*, edited by Milton K. Munitz, 199–211. New York University Press.

MacFarquhar, L. 2025. “Some People Can’t See Mental Images. The Consequences Are Profound.” *The New Yorker*, November 3, 2025.

Manning, L. 2015. “No Identity without an Entity.” *Pacific Philosophical Quarterly* 96(2): 279–305.

Martin, M. G. F. 1997. “The Reality of Appearances.” In *Thought and Ontology*, edited by M. Sainsbury, 81–106. FrancoAngeli.

———. 2002a. “Particular Thoughts and Singular Thought.” *Royal Institute of Philosophy Supplement* 51: 173–214.

———. 2002b. “The Transparency of Experience.” *Mind & Language* 17(4): 376–425.

———. 2004. “The Limits of Self-Awareness.” *Philosophical Studies* 120: 37–89.

———. 2006. “On Being Alienated.” In *Perceptual Experience*, edited by T. S. Gendler and J. Hawthorne, 354–410. Oxford University Press.

———. 2010. “What’s in a Look?” In *Perceiving the World*, edited by B. Nanay, 160–225. Oxford University Press.

———. 2020. “The Diversity of Experiences.” *Philosophy and Phenomenological Research* 100(3): 728–37.

Matravers, D. 1997. “The Paradox of Fiction: The Report versus the Perceptual Model.” In *Emotions and the Arts*, edited by M. Hjort and S. Laver, 78–92. Oxford University Press.

McDowell, J. 1977. “On the Sense and Reference of a Proper Name.” *Mind* 86(342): 159–85.

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

———. 1982b. “Truth-Value Gaps.” *Studies in Logic and the Foundations of Mathematics* 104: 299–313.

———. 1987. “Singular Thought and the Extent of Inner Space.” In *Subject, Thought and Context*, edited by P. Pettit and J. McDowell, 137–168. Clarendon Press.

———. 1994. *Mind and World*. Harvard University Press.

Meinong, A. (1904) 1960. “The Theory of Objects.” Translated by I. Levi, D. B. Terrell, and R. M. Chisholm. In *Realism and the Background of Phenomenology*, edited by R. M. Chisholm, 76–117. The Free Press.

Montague, R. 1973. “The Proper Treatment of Quantification in Ordinary English.” In *Approaches to Natural Language*, edited by Jaakko Hintikka, Julius M. E. Moravcsik, and Patrick Suppes, 221–42. D. Reidel.

Moore, G. E. 1905–6. “The Nature and Reality of Objects of Perception.” *Proceedings of the Aristotelian Society* 6: 68–127.

———. 1918–1919. “Some Judgments of Perception.” *Proceedings of the Aristotelian Society*, n.s., 19: 1–29.

———. 1925. “A Defence of Common Sense.” In *Contemporary British Philosophy*, 2nd series, edited by J. H. Muirhead, 192–233. George Allen & Unwin.

———. 1952. “A Reply to My Critics.” In *The Philosophy of G. E. Moore*, edited by P. Schilpp, 533–667. Northwestern University Press.

Moran, R. 1994. “The Expression of Feeling in Imagination.” *Philosophical Review* 103(1): 75–106.

Neale, S. 1988. “Events and ‘Logical Form.’” *Linguistics and Philosophy* 11(3): 303–21.

Neill, A. 1991. “Fear, Fiction and Make-Believe.” *Journal of Aesthetics and Art Criticism* 49: 47–56.

Nelkin, N. 1989. “Consciousness and the Origins of Thought.” *Philosophical Studies* 56: 421–39.

Noonan, H. W. 1991. “Object-Dependent Thoughts and Psychological Redundancy.” *Analysis* 51(1): 1–9.

O’Brien, L. 2017. “Actions as Prime.” *Royal Institute of Philosophy Supplement* 80: 265–85.

Parsons, T. 1990. *Events in the Semantics of English*. MIT Press.

Pautz, A. 2007. “Intentionalism and Perceptual Presence.” *Philosophical Perspectives* 21: 495–541.

Peacocke, C. A. B. 1983. *Sense and Content*. Clarendon Press.

Perry, J. 1980. “A Problem About Continued Belief.” *Pacific Philosophical Quarterly* 61: 317–22.

———. 2001. *Reference and Reflexivity*. CSLI Publications.

Pitcher, G. 1971. *A Theory of Perception*. Princeton University Press.

Pitt, D. 2004. “The Phenomenology of Cognition, Or, What Is It Like to Think That P?” *Philosophy and Phenomenological Research* 69(1): 1–36.

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

———. 1940. “Some Philosophical Questions about Telepathy and Clairvoyance.” *Philosophy* 15: 363–74.

Priest, G. 2005. *Towards Non-Being: The Logic and Metaphysics of Intentionality*. Clarendon Press.

Prinz, J. J. 2004. *Gut Reactions: A Perceptual Theory of Emotion*. Oxford University Press.

Prior, A. N. 1971. *Objects of Thought*. Edited by P. T. Geach and A. J. P. Kenny. Oxford University Press.

Pylyshyn, Z. 1989. "The Role of Location Indexes in Spatial Perception: A Sketch of the FINST Spatial-Index Model." *Cognition* 32(1): 65–97.

Quine, W. V. O. 1956. "Quantifiers and Propositional Attitudes." *Journal of Philosophy* 53: 177–87.

———. 1960. *Word and Object*. MIT Press.

———. 1969. *Ontological Relativity and Other Essays*. Columbia University Press.

Recanati, F. 2012. *Mental Files*. Oxford University Press.

Reid, T. 1785. *Essays on the Intellectual Powers of Man*. J. Bell and G.G.J. & J. Robinson.

Robinson, H. 1994. *Perception*. Routledge.

Rorty, R. 1979. *Philosophy and the Mirror of Nature*. Princeton University Press.

Russell, B. 1905. "On Denoting." *Mind* 14(56): 479–93.

———. 1906. "On the Nature of Truth." *Proceedings of the Aristotelian Society* 7: 28–49.

———. (1910) 1992. "On the Nature of Truth and Falsehood." In *The Collected Papers of Bertrand Russell*, vol. 6, *Logical and Philosophical Papers, 1909–13*, edited by J. G. Slater and B. Frohmann, 116–124. Routledge.

———. 1910–1911. "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society* 11: 108–28.

———. 1912. *The Problems of Philosophy*. Galaxy.

———. 1913. *Theory of Knowledge: The 1913 Manuscript*. Edited by E. R. Eames and K. Blackwell. Routledge.

———. (1918) 1986. "On Sensations and Ideas." In *The Collected Papers of Bertrand Russell*, vol. 8, *The Philosophy of Logical Atomism and Other Essays, 1914–19*, edited by J. G. Slater, 252–55. George Allen & Unwin.

Sainsbury, R. M. 2000. "Evans on Reference." In *Departing from Frege*, 43–70. Routledge.

———. 2005. *Reference without Referents*. Oxford University Press.

———. 2010. "Intentionality without Exotica." In *New Essays on Singular Thought*, edited by Robin Jeshion, 300–18. Oxford University Press.

———. 2018. *Thinking About Things*. Oxford University Press.

———. 2020. “Varieties of Singularity.” In *Singular Thought and Mental Files*, edited by R. Goodman et al., 21-37. Oxford University Press.

Salmon, N. 1986. *Frege’s Puzzle*. MIT Press.

———. 1998. “Nonexistence.” *Noûs* 32(3): 277–319.

———. 2002. “Mythical Objects.” In *Meaning and Truth: Investigations in Philosophical Semantics*, edited by Joseph Keim Campbell, Michael O’Rourke, and David Shier, 105–123. Seven Bridges Press. Reprinted in *Metaphysics, Mathematics, and Meaning: Philosophical Papers, Volume I*, 91–108. Clarendon Press, 2005.

Scholl, B. J., and Z. W. Pylyshyn. 1999. “Tracking Multiple Items Through Occlusion: Clues to Visual Objecthood.” *Cognitive Psychology* 38: 259–90.

Searle, J. 1969. *Speech Acts*. Cambridge University Press.

———. 1983. *Intentionality*. Cambridge University Press.

Segal, G. 1989. “The Return of the Individual.” *Mind* 98(389): 39–57.

Sellars, W. 1956. “Empiricism and the Philosophy of Mind.” *Minnesota Studies in the Philosophy of Science* 1: 253–329.

Sethi, U. 2020. “Sensible Over-Determination.” *Philosophical Quarterly* 70(280): 588–616.

Siegel, S. 2010. *The Contents of Visual Experience*. Oxford University Press.

Siewert, C. 1998. *The Significance of Consciousness*. Princeton University Press.

Smythies, J. R. 1956. *Analysis of Perception*. Routledge.

Soames, S. 1987. “Direct Reference, Propositional Attitudes, and Semantic Content.” *Philosophical Topics* 15(1): 47–87.

———. 2002. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford University Press.

———. 2003. *Philosophical Analysis in the Twentieth Century, Vol. 2: The Age of Meaning*. Princeton University Press.

Stalnaker, R. 2009. “What Is De Re Belief?” In *The Philosophy of David Kaplan*, edited by J. Almog and P. Leonardi, 233–45. Oxford University Press.

Stocker, M. 1987. “Emotional Thoughts.” *American Philosophical Quarterly* 24: 59–69.

Strawson, P. F. 1959. *Individuals*. Methuen.

———. 1974. *Subject and Predicate in Logic and Grammar*. Methuen.

Stroud, B. 1968. “Transcendental Arguments.” *Journal of Philosophy* 65(9): 241–56.

———. 2009. “Scepticism and the Senses.” *European Journal of Philosophy* 17(4): 559–70.

Thomason, R. H., and R. C. Stalnaker. 1973. “A Semantic Theory of Adverbs.” *Linguistic Inquiry* 4(2): 195–220.

Tieszen, R. 2002. “Gödel and the Intuition of Concepts.” *Synthèse* 133(3): 363–91.

Tiskin, D. 2018. “Intentional Identity as a Transparency Phenomenon.” In *Pronouns in Embedded Contexts at the Syntax-Semantics Interface*, edited by Pritty Patel-Grosz, Patrick Georg Grosz, and Sarah Zobel, 43–73. Springer.

Travis, C. 2004. “The Silence of the Senses.” *Mind* 113: 57–94.

Tye, M. 1992. “Visual Qualia and Visual Content.” In *The Contents of Experience*, edited by T. Crane, 158–176. Cambridge University Press.

———. 1995. *Ten Problems of Consciousness*. MIT Press.

u/lionellrichie. 2020. “Just Realized My Soap Wasn’t Working Because It’s Literally a Block of Cheese.” Reddit, r/funny, March 9, 2020. https://www.reddit.com/r/funny/comments/fg58qo/.

van Hulst, C., and M. Cresswell, eds. 2016. “A. N. Prior on Austin’s Sense and Sensibilia.” *Journal for the History of Analytical Philosophy* 4(4): 1–41.

Viberg, Å. 1984. “The Verbs of Perception: A Typological Study.” In *Explanations for Language Universals*, edited by B. Butterworth et al., 123–62. de Gruyter.

Walton, K. 1973. “Pictures and Make-Believe.” *Philosophical Review* 82(3): 283–319.

———. 1978. “Fearing Fictions.” *Journal of Philosophy* 75(1): 5–27.

———. 1990. *Mimesis as Make-Believe*. Harvard University Press.

———. 2015. “Empathy, Imagination, and Phenomenal Concepts.” In *In Other Shoes: Music, Metaphor, Empathy, Existence*, 1–16. Oxford University Press.

White, Alan R. 1958. *G. E. Moore: A Critical Exposition.* Oxford: Basil Blackwell.

Whitehead, A. N., and B. Russell. 1910. *Principia Mathematica*. Vol. 1. Cambridge University Press.

Wiggins, D. 1995. “The Kant–Frege–Russell View of Existence: Toward the Rehabilitation of the Second-Level View.” In *Modality, Morality, and Belief: Essays in Honor of Ruth Barcan Marcus*, edited by W. Sinnott-Armstrong, D. Raffman, and N. Asher, 41–67. Cambridge University Press.

Williams, Bernard. (1978) 2005. *Descartes: The Project of Pure Enquiry*. Routledge.

Williamson, T. 1995. “Is Knowing a State of Mind?” *Mind* 104(415): 533–65.

———. 1998. “The Broadness of the Mental.” *Philosophical Perspectives* 12: 389–410.

———. 2002. *Knowledge and its Limits*. Oxford University Press.

Wilson, G. 1989. *The Intentionality of Human Action*. Stanford University Press.

Wittgenstein, L. 1922. *Tractatus Logico-Philosophicus*. Translated by F. Ramsey and C. Ogden. Harcourt, Brace & Company.

———. 1953. *Philosophical Investigations*. Translated by G. E. M. Anscombe. Blackwell.

Zeman, A., F. Milton, S. Della Sala, M. Dewar, T. Frayling, J. Gaddum, A. Hattersley, et al. 2020. "Phantasia—The Psychological Significance of Lifelong Visual Imagery Vividness Extremes." *Cortex* 130: 426–40.