2
Folk Psychology and Its Ontology

2.1 Introduction

When we ask about imagination's relation to states like beliefs, desires, intentions, judgments, decisions, and so on, our answers will turn, in part, on what we take those other states to be—on how we view *their* ontological status. To inquire after the ontological status of such states is to ask for a deeper account of their nature than the platitudinous definitions we might find in a dictionary. Two people can agree that to believe something is to take it to be true, for instance, while disagreeing about what beliefs *are* in a deeper sense. Likewise, they may agree that to imagine is to engage in rich, epistemically blameless thought about the possible, fantastical, or fictional while disagreeing about the deeper nature of imaginings. Notoriously, there are rather different views in philosophy concerning the "deep" nature of folk psychological states. This chapter surveys some of those views with an eye toward explaining how their differences bear on the project of explaining imagination.

One goal is to show that the project of explaining imagination, as pursued here, is open to researchers who don't share assumptions about the nature of folk psychological states. For instance, many debates about imagination occur among theorists who share a background belief in the representational theory of mind (Aydede, 2015; Fodor, 1987; Nichols, 2006a). But one needn't accept that theory of mental representation and its relation to folk psychology in order to find the project of explaining imagination both approachable and important. I will argue that, whatever your take on folk psychological ontology may be, a theory that breaks imagination into states like beliefs, desires, and intentions has the potential to offer a genuine explanation of imagination. This is so even if you are an *eliminativist* about folk psychological states (Churchland, 1981), or even if you think that cognitive science has no need for the notion of *mental representation* (Chemero, 2011).

My second goal in this chapter is to defend the reductive style of explanation I pursue against a few objections. These are not objections to specific examples or contexts where I propose that imagination-talk can be replaced with talk of other kinds of states. They are instead objections to the effect that, *even if* we could replace imagination-talk with talk of other kinds of folk psychological states—in ways I previewed last chapter—this still would not constitute an explanation (or *reduction*) of imagination. This sort of objection is best addressed by distinguishing different views one might have on folk psychological ontology; that is why I take it up here, in sections 2.6 and 2.7.

Explaining Imagination. Peter Langland-Hassan, Oxford University Press (2020). © Peter Langland-Hassan.
DOI: 10.1093/oso/9780198815068.001.0001

A last goal of this chapter is to clarify how different sources of evidence bear on the question of explaining imagination. In making claims about imagination and its relationship to other mental states, we seem to pronounce on the structure and nature of a biological phenomenon: the human mind. It is not always clear how such claims could be supported by the kinds of considerations philosophers typically adduce—deriving, for instance, from introspection, the mining and refinement of commonsense platitudes, and appeals to parsimony. (We are well advised by Bechtel & Richardson (2010) to expect complexity, not parsimony, when discovering the nature of organisms.) On the other hand, it can also be hard to see how a harder-nosed empirical approach could gain better traction on the questions that concern us. Clarity on these matters comes when we recognize that the proper epistemological approach to explaining imagination will depend in part on our broader views concerning folk psychological ontology, in ways I hope to elucidate.

As this chapter is largely meta-theoretical in nature, it can be skipped without compromising one's ability to follow most of the arguments in later chapters. There will, however, be places later on where the distinctions drawn here—between "heavy-duty" and "light-duty" ontologies, for example—are essential to grasping the issues at play. Also, for any who wondered, last chapter, whether the kinds of explanations I'll pursue are *explanations* in good standing, this chapter is essential reading. So the recommended approach is to take time now to draw the distinctions we'll need later, and to confirm that we're on solid explanatory footing.

2.2 Folk Psychological Ontologies—a Brief History

One of the great innovations of twentieth-century philosophy was the idea that everyday psychological terms—words like 'belief,' 'desire,' and 'imagining'—could be seen as *theoretical* terms (Sellars, 1956). The supposed theory featuring those terms came to be called *folk psychology*. According to legend, this was the theory of the *folk*—ordinary folk you might see at the post office, or waiting in line to vote. Not that they would have told you they had a psychological theory. But, stepping back, we could view them as using one—one that enabled them to understand and predict others' behavior by attributing to them states like beliefs, desires, and intentions.

Why was Jason taking off his shoes? The folk could explain: he *desired* to pass through airport security and *believed* he must remove his shoes to do so. Why wasn't Jim? Well, he didn't *believe* it was required. Why was Julia leaping over hurdles? She *desired* to win the race and *believed* that jumping the hurdles would be faster than running through them. Why was Julia's mom so happy? She desired that Julia would win and believed that Julia was winning.

The philosophical attraction in this was that we could avoid relying upon introspection to identify and categorize mental phenomena, instead treating mental states as unobservable entities that, like electrons or quarks, are *posited* in

order to explain phenomena that we can all observe together (Chihara & Fodor, 1965). Psychological notions like belief and desire could earn their keep in the same way that other theoretical entities do: by their usefulness to explanations of outwardly observable phenomena. And the best part of it was that, despite our having done no formal experiments, we already had the relevant theory in hand—"folk" psychology—just as a function of being competent speakers of a natural language that incorporates mental state terms like 'belief' and 'desire.'

It would be hard to overstate the impact this doctrine has had on how philosophy and psychology approach the study of mental states and processes. Within philosophy, in particular, there arose near universal agreement that folk psychology is a powerful and useful means for predicting and explaining human behavior. It seems we would be hard-pressed to say *why* Jason is removing his shoes if we were *not* allowed use of any folk psychological terms. Simply appealing to environmental context—saying that Jason is removing his shoes *because it is required*—doesn't explain why Jim, who is standing next to him, fails to do so. The problem is solved if we can reference their respective states of mind; and attributing different folk psychological states is a good way of doing that. Sure, there might be some other story to tell that would distinguish Jason from Jim—one involving retinal stimulation, neural firings, and the like. Be we're not yet in any position to tell it. And who's to say it would offer a *better* explanation? For power and ease of use, folk psychology is hard to beat.

And yet, despite wide agreement about the power and usefulness of folk psychology *as a practice* for predicting, explaining, and rationalizing behavior, there has never been a consensus concerning the ontology it implies. Exactly what *sort* of things (if any) are we claiming to exist when we grant that folk psychological terms offer useful means for predicting and explaining behavior? Answering requires us to clarify the distinction between folk psychological *talk*—that is, our everyday practice of attributing folk psychological states like beliefs, desires, and imaginings, to each other—and the actual states of our minds, brains, bodies, and environments that are causally responsible for our behavior. For some in contemporary philosophy—including many who work on imagination—there is little distance between the two. According to this family of theorists, to say that Jane believes that p is just to say that Jane has, realized in her brain, a *mental representation* of a certain sort—one with p as its content. This mental representation is then thought to play various causal roles in shaping her behavior—behavior on the basis of which we infer that she believes that p. It is thought that we can move from the truth of a folk psychological description—that Jane believes that p—to a specific (albeit defeasible) claim about the structure of Jane's mind—viz., that it contains a mental representation with the content p (Dretske, 1991; Fodor, 1987).

It is important to see why this inference, correct or not, is far from inevitable. Not all who find folk psychological *talk* explanatorily useful feel obliged to posit corresponding mental representations. A second family of theories finds it useful to attribute folk psychological states, but resists any move from there to the

conclusion that the mind contains mental representations of a particular sort. Some in this family hold that cognitive science will not include folk psychological states in its account of what the mind contains, and, on those grounds, advocate eliminativism about the states (Churchland, 1981; Stich, 1983). Others take a dispositionalist view of folk psychological states, holding that to have a certain belief or desire is simply to fulfill a certain dispositional stereotype (Ryle, 1949/2009; Schwitzgebel, 2002, 2013). This approach makes no comment on the causal bases of those dispositions and, typically, views with skepticism the idea the causal bases are mental representations whose contents mirror the contents of the that-clauses featured in folk psychological talk. (The "that-clause" in the folk psychological ascription, "Jim believes that Mars is hot," is "Mars is hot"). Also in this family is Daniel Dennett, who holds that to have folk psychological states is simply to be the sort of "intentional system" whose behavior can be explained and predicted by attributing to it such states (with their associated dispositions)—again without comment on the nature of the internal features of the system that make it suitable for description in such terms (Dennett, 1989, 1991). Still others in this group defend a "minimalist" approach, holding that folk psychological states are semantically evaluable, causally efficacious internal states, while eschewing any commitments about whether this implies the existence of corresponding mental representations with a semantics that roughly matches that of the that-clauses used in ordinary folk psychological attributions (Egan, 1995; Graham & Horgan, 1988).

In short, two broad families of theory—each with influential members—agree that it is useful to ascribe beliefs, desires, imaginings, and the like to people when predicting and explaining their behavior. But they disagree on the sort of things that are being ascribed when we say of someone that she believes or desires that *p*. It will be useful to look more closely at each approach now in order to appreciate how their differences bear on the project of explaining imagination.

2.3 Heavy-Duty Ontology

The most general commitment uniting the first family of theories—what I will call *heavy-duty* views of folk psychological ontology—is that folk psychological mental state ascriptions refer to discrete mental representations tokened in individuals, where the semantics (or meaning) of these representations typically bears a close relationship to the semantics of the that-clauses we use to ascribe them. On this view, when we say that Joe believes (or desires, or intends) that *there is coffee in the mug*, the statement is made true by the fact that Joe has a mental representation realized in his brain with the content *there is coffee in the mug* (or something semantically close to that)—where this mental representation has the distinctive causal role of a belief (or desire, or intention). It is the causal

interaction of such mental representations with each other that, on the heavy-duty view, serves to bring about the behaviors or dispositions that we predict and explain through folk psychological talk. So characterized, the heavy-duty view is a close cousin to the representational theory of mind (RTM) in philosophy and psychology; the two views only diverge if defenders of RTM don't insist on a close relation between the semantics of (at least many of) the mental representations used in human cognition and semantics of the that-clauses typically used in folk psychological state attributions.[1] Others have called this sort of view *intentional realism* (Pitt, 2020). That label strikes me as pejorative, however, as it wrongly implies that rejecting it makes one an anti-realist about folk psychological states (more on this later).

The most famous heavy-duty view comes from Jerry Fodor—especially Fodor (1975) and (1987, Ch. 1). For Fodor, it is just because our internal mental representations closely mirror—in both their syntactic structure and semantics—the natural language sentences we use to describe someone in folk psychological terms that our commonsense view of ourselves as rational agents stands to be vindicated. The idea that we act *for reasons*—reasons we are able to describe ourselves as having—can be seen to cohere with our being causally efficacious parts of the physical world, he argues, if the causes of our behaviors are internal representations that share semantic properties with (relevant portions of) the folk psychological sentences we apply to ourselves. One of the key thoughts inspiring Fodor and his followers is that, with the development of computers, it becomes possible to see beliefs and desires *both* as having meanings *and* as being physical states in the brain. The analogy of thinking to computing allows us to see how it is possible for a system to be set up so that the causal interactions that occur among its internal states (as a function of their intrinsic physical properties or "shape") mirror the inferential relationships we would expect to hold among symbols with certain meanings. Patterns of semantic entailment—sentence A rationally entailing sentence B—are realized in sequences of physical symbols whose causes and effects "contrive to respect" the semantic values we've assigned to them (Fodor, 1987, pp. 10–20; Aydede, 2015).

An important feature of the Fodorian version of the heavy-duty view is that mental representations have a relational structure, involving a mental sentence—one with a particular *meaning* or *content*—and an attitude taken toward that sentence. On most iterations of this view, the "mental sentences" in our heads don't

[1] Typically, defenders of RTM posit mental representations whose semantics *do* closely mirror the semantics of ordinary folk psychological state attributions. However, there is room in logical space for someone to defend a representational theory of mind without holding that the mental representations used in human cognition bear an appreciable relation to those of the sentences we use to attribute folk psychological states. This is why I have defined *heavy-duty* views so as to explicitly require a close mirroring between the semantics of mental representations and those of the that-clauses used in ordinary folk psychological state attributions.

occur in a natural, spoken language, but rather in a proprietary "language of thought," sometimes called *Mentalese* (Fodor, 1975). Not only do these representations have meanings that closely mirror the meanings of the that-clauses used in folk psychological attributions, they also are said to have a language-like syntactic structure, insofar as they are composed of discrete meaningful symbols, where the meaning of a complex representation (e.g., a belief) is a function of the meaning of its parts, together with the syntactic rules for combining them.

Whether a mental representation qualifies as a belief, desire, or some other kind of state is then said to be determined by the causal-functional role of the representation in the broader cognitive economy. Bearing the relation of belief, as opposed to desire, to a mental representation with the content *p* will be a matter of the kinds of causes and effects the state has—its "functional role." I will call the different causal-functional profiles characteristic of different kinds of folk psychological states *psychological attitudes*. So, where ordinary folk psychology speaks of *believing* that *p*, *desiring* that *q*, *wondering whether r*, and so on—these being different "attitudes" one can take toward the propositions *p*, *q*, and *r*—the heavy-duty theorist posits corresponding *psychological* attitudes that are different relations one can bear to mental representations with the contents *p*, *q*, or *r*. Unlike the notion of a (mere) propositional attitude, the notion of psychological attitude is intended to carry with it the idea that there are mental representations tokened in one's brain toward which one takes the relevant attitude, where one's taking the attitude is to be understood in terms of the representation's having a certain functional role in one's cognitive economy.

Often, theories that posit psychological attitudes follow Schiffer (1981) and Fodor (1987) in speaking of "boxes" corresponding to each attitude; these boxes are meant to summarize, within a diagram, the kinds of causes and effects distinctive of each attitude-type. So, to believe that *p* is to have a representation with the content *p* "in" one's Belief Box and to desire that *q* is to have a representation with *q* "in" one's Desire Box. The boxes are not assumed to have any geographic reality in the mind itself; boxes, *qua* boxes, exist only in the diagrams meant to map out the causal-functional relations among mental representations with different contents. The use of the box metaphor *does*, however, presume the existence of certain kinds of mental representations that reside "in" the boxes, insofar as those representations have certain causes and effects. Specifically, it assumes mental representations whose contents (or semantics) closely mirror those of the that-clauses we would use to accurately describe someone in folk psychological terms. Note that this does not require any further assumption that the representations are language-like in structure. So, while many heavy-duty views come with specific commitments about the format of the mental representations in the boxes they posit (viz., that they are language-like), the only commitment I attribute to all heavy-duty theorists is the idea that the success of our folk psychological talk is, in general, explained by the existence of mental representations with a closely

matching *semantics*—where each representation has a discrete location in space and time. Further, heavy-duty theorists needn't hold that *every* instance of a successful folk psychological explanation is itself explained by the presence of semantically-matching mental representations (hence the "in general" above). They can allow that a formal cognitive scientific inventory of one's mental states will involve some "cleaning up"—or even dismissal of—ordinary folk psychological talk. They are simply committed to the final inventory including representations whose semantics have a fairly transparent relationship to the semantics of the that-clauses we use in ordinary folk psychological explanations, and to the idea that such representations account for the usefulness of folk psychological talk *most of the time*.

I have found these claims about what "boxes" presuppose to be controversial in some quarters. It has been objected to me that box-talk is simply shorthand for functionalism in general, and needn't commit one to the existence of mental representations of any sort. I think that is incorrect. Within cognitive psychology, box-and-arrow diagrams are intended to map the flow of information through the mind and brain. Typically, a diagrammatic distinction is made between boxes, which represent data stores, and hexagons, which represent mechanisms capable of operating on the data stores (see, e.g., Nichols & Stich, 2000, p. 121). The distinction between a data store, on the one hand, and a mechanism that operates on the data, on the other, is at odds with a "merely functionalist" picture, where mental states are defined in terms of their functional roles, without comment on corresponding mental representations—representations that have discrete locations in time and space. After all, for a mechanism to operate on a mental state, the state must be physically realized in some form; one's being in the state cannot simply be a matter of one's having certain dispositions (as on some of the "light-duty" functionalist views discussed below). So talk of boxes and mechanisms thus brings with it the need for mental representations that are tokened "in" the boxes, such that other mental mechanisms can transform them in various ways. I will assume as much going forward in my use of "box" terminology. (For a functionalist picture of folk psychological ontology that lacks any commitment to corresponding mental representations, see Egan (1995).)

In recent decades, many have proposed that imagination involves use of a proprietary psychological attitude as well—one with similarities to belief, but which is ultimately quite distinct (see, e.g., Carruthers, 2006, pp. 89–91; Currie & Ravenscroft, 2002, Ch. 2; Friedman & Leslie, 2007, p. 115; Gendler, 2006, pp. 183–5; Nichols & Stich, 2003; Schellenberg, 2013; Schroeder & Matheson, 2006; Spaulding, 2015; Stokes, 2014; Weinberg & Meskin, 2006b). In some cases, they go so far as to posit an "Imagination Box" (Doggett & Egan, 2007; Liao & Doggett, 2014; Nichols, 2008; Schellenberg, 2013; Weinberg & Meskin, 2006b). This view is at odds with the reductive account I will pursue.

2.4 Light-Duty Ontology

In characterizing heavy-duty views, I highlighted a distinction between what is a quasi-scientific hypothesis about the nature of our minds—viz., that they contain mental representations with specific contents and functional roles—and what is something that anyone who successfully makes use of folk psychological descriptions must grasp. The latter includes more superficial phenomena, *behavioral dispositions* central among them. The competent user of the term 'belief' understands that someone who believes that *p* and desires that not-*p* has certain characteristic dispositions to behavior, whether or not they have any views about what it is that gives the person those dispositions—just as one might know that a vase is fragile, and so disposed to break when dropped, without having any clear idea of what it is about the vase that makes it fragile.

Folk psychology aside, we routinely ascribe dispositions to people on the basis of noticing superficial features that are reliable markers for the dispositions, without any understanding of the causal bases for the dispositions. Noticing that a husband and wife are both tall and blonde, we infer that they are disposed to have tall, blonde children. We needn't have any idea of the causal bases (grounded in their genetics) for those dispositions, in order to exploit knowledge of the dispositions in making predictions about their offspring. We move from superficial features we can observe, to knowledge of associated dispositions, to predictions and explanations of specific phenomena. In the same way, what I will call the *light-duty view* holds that we are able to infer, on the basis of a person's superficial behavior (and context), dispositions they are likely to have. Our folk psychological ascriptions, made on the basis of observed behavior, serve to attribute dispositions that will further manifest in their future behavior. Thus we can predict and explain specific behaviors on the basis of their having the dispositions we ascribe with the use of folk psychological terms—even if we remain clueless about the causal bases for the dispositions. Light-duty views take these superficial phenomena to capture the essence of folk psychological states.

To get a better grasp on this, consider David Lewis's (1972) distinction between the causal-functional role of a mental state and the occupant of that role. The causal role of a mental state, Lewis held, can be extracted from the set of platitudes that competent speakers of the language accept about the state. These "roles" are dispositional in nature. Lewis characterizes them thus:

> When someone is in so-and-so combination of mental states and receives stimuli of so-and-so kind, he tends with so-and-so probability to be caused thereby to go into so-and-so mental states and produce so-and-so motor responses.
>
> (Lewis, 1972, p. 256)

The dispositions Lewis lists are both dispositions to have certain behavioral ("motor") responses and dispositions to go into other mental states. These mental

states themselves *could* be understood as mental representations, in the manner of the heavy-duty view; but they can also be understood, more superficially, as *states of having certain further dispositions*, without comment on the causal bases of the dispositions. This more cautious, superficial understanding of folk psychological states is where light-duty views set up shop (see, e.g., Schwitzgebel, 2002, 2013).

Lewis himself thinks of mental states as the "occupants" of the causal-functional roles we extract from folk psychological platitudes: "When we learn what sort of states occupy those causal roles definitive of the mental states," he writes, "we will learn what the mental states are… exactly as we found out what light was when we found that electromagnetic radiation was the phenomenon that occupied a certain role" (1972, p. 256). The key difference between Lewis and heavy-duty theorists, as characterized above, is that he assumes we do not *yet* know what the mental states are. We just know that, if they exist, they will be the occupants of certain causal roles; they will be the states that cause people to have the dispositions we attribute to them when we attribute them beliefs, desires, and the like. By Lewis's lights, these occupants *might* be mental representations of certain kinds; or they might be non-representational neurobiological states; or they might—with less likelihood—be conglomerations of glue and sawdust. Our expertise with folk psychological explanation does not prejudge an answer (though our broader understanding of nature and biology might). The heavy-duty theorist, by contrast, has in mind an account of what those occupants are: mental representations, realized in the brain, with contents mirroring those of the that-clauses used in appropriate folk psychological descriptions.

So both light- and heavy-duty views will agree that if Joe *desires* to keep a dying fire lit and *believes* that adding another log will do the trick, then, all else equal,[2] he will add another log. In ascribing such a belief and desire pair to Joe, both heavy and light views agree that Joe has a number of interesting dispositions, such as to agree with others that the fire should be kept lit, to assist in searching for a log, to be pleased when the fire remains lit, and so on. The light-duty conception remains "light" in making no comment on the nature of the internal states in virtue of which Joe has those dispositions; whereas, on the heavy-duty view, when Joe believes that the fire is almost out, there is a representation realized in Joe's brain whose meaning is that the fire is almost out; this representation causally interacts with other mental representations so as to *result in* his having log-adding dispositions.

Whether folk psychological ascriptions are ever strictly speaking *true* is answered in different ways by different light-duty theorists. Eliminativists hold that the ascriptions are strictly speaking *false*, despite their frequent utility

[2] The *all else equal* clause is notoriously difficult to fill in. To start, Joe must not have a stronger countervailing desire; he must not believe there is a better, easier, way to keep the fire lit; he must believe he is allowed to add a log; and so on. These difficulties are shared by both the light- and heavy-duty views.

(Churchland, 1981). Other light-duty theorists remain agnostic concerning the truth or falsity of the ascriptions. For instance, a light-duty theorist may, like Lewis, identify folk psychological states themselves with their causal bases, *whatever they turn out to be*. (Lewis leaves open the possibility that there will be no unified realization base for the causal-functional roles and, in that case, appears ready to conclude that no such states exist (Lewis, 1972, p. 252).) Alternatively, a light-duty theorist may identify being in a folk psychological state simply with the possession of certain dispositions, and not with any putative causal bases for the dispositions (Schwitzgebel, 2002, 2013; Ryle, 1949; Sellars, 1956). For instance, on Eric Schwitzgebel's "phenomenal dispositionalist" view, being in a certain folk psychological state amounts to "having a dispositional profile that matches, to an appropriate degree and in appropriate respects, a stereotype for that attitude, typically grounded in folk psychology" (Schwitzgebel, 2013). (He includes within such dispositional profiles "phenomenal dispositions" to have certain kinds of conscious experiences (Schwitzgebel, 2002, p. 252).) Schwitzgebel contrasts his "superficial," dispositional account of the attitudes to "deep" views of the Fodorian kind. And, indeed, Schwitzgebel's distinction between "superficial" and "deep" views of folk psychological states aligns closely with my distinction between "light-duty" and "heavy-duty" views. (I'm indebted to Schwitzgebel's description of the terrain, though I don't wish to saddle him with my slightly different understanding of it.)

While each person who believes that *p* will have dispositions in common with every other person who believes that *p*—provided their *other* relevant folk psychological states are similar enough—there is, on the light-duty view, no expectation that we will find an interesting *type* of internal state shared by all and only those who believe that *p*—one that makes it the case that they have those dispositions. In individual cases, we may be able to answer the question: what is it about *S* that makes him have the dispositions associated with believing that *p*? But, broadening our search for the *more general* internal causes of the dispositions we associate with believing that *p*, we may find only a messy disjunction of different kinds of states. Light-duty theorists, including Schwitzgebel, Dennett (1991), and Egan (1995), are typically skeptical that cognitive science will discover mental representations realized in the brain with contents mirroring the meanings of the that-clauses used in folk psychological ascriptions. While they can leave the door open to such a discovery, their hunch is that folk psychological notions like *believing that p* will break into many different neuro-cognitive pieces when it comes to discovering their implementation in individual systems.

Dominic Murphy gives voice to this view in a paper on the place of folk psychology in cognitive science:

> The question whether science makes use of representational systems isn't really open to doubt any longer: many areas of psychology and neuroscience take for

> granted the existence of semantically interpretable internal states… What is open to doubt is whether representation, as used in the sciences of the mind, has the properties that philosophers have found in intentional content, as presupposed by folk psychology. (Murphy, 2017, p. 138)

Murphy goes on to articulate a light-duty view that still finds an important role for folk psychological notions in cognitive science:

> The concept of belief will do very little useful explanatory work in any mature cognitive science. But it might nevertheless be decomposable into a family of successor notions that can suggest and guide useful neuroscientific hypotheses. (p. 138)

Note that, while Murphy thinks that belief will not be a central notion in a mature cognitive science, he suspects it will play an important role as a kind of *ancestor* notion, the exploration, refinement, and revision of which will constitute crucial steps in understanding how the mind really works. For that reason, the notion retains value in the here-and-now.

Most light-duty theorists, like Murphy, will allow that there are mental representations of some sort underlying human cognition; they just doubt that the contents (or semantics) of those representations bear any appreciable relation to the contents of the that-clauses featured in folk psychological ascriptions. Nevertheless, they *need not* hold that there are any such mental representations in order to maintain that folk psychological ascriptions are true—true either because one's having certain dispositions suffices for their truth (as in Schwitzgebel's view), or because being in any kind of internal state at all that leads one to have those dispositions—no matter how disjunctive it may be across cases—suffices for their truth. At the limit, a light-duty theorist can hold that folk psychological ascriptions are for the most part true, while maintaining that a mature cognitive science will have no use for the notion of mental representation at all (Chemero, 2011).

2.5 Heavy-Duty Incredulity about Light-Duty Dispositionalism, and Principled Agnosticism

Those with heavy-duty views sometimes react to the light-duty perspective with incredulity. How, they ask, does the light-duty theorist propose to *explain* all the dispositions we cite so regularly, *other than* by positing internal representations of a heavy-duty sort (Fodor, 1987; Quilty-Dunn & Mandelbaum, 2017)? This sort of incredulity is worth discussion, as it helps to clarify what is at stake in debates about folk psychological ontology.

Light-duty theorists can push back in several ways. First, they can hold that there are, in fact, other well-developed possibilities for explaining the

dispositions. For instance, a light-duty theorist may think that *connectionist networks* offer a better model for how the brain accomplishes the information-processing relevant to explaining human behavior (P. S. Churchland & Sejnowski, 1989; Van Gelder, 1990, 1998). Such networks are standardly held to involve mental representations and computations over those representations. Yet these representations don't have contents that mirror those of the sentences we use to describe someone in folk psychological terms. In a network set up to identify images of dogs, for instance, there are no representations with the content "dogs have four legs," or "dogs have hair." Instead, the networks have characteristic patterns of activation, according to the "weights" assigned to different connected nodes in the network (where the connection weights between nodes are intended to mirror the connection strengths between neurons, or sets of neurons). Whatever semantic relationships hold among different states of these networks—in virtue of which they qualify as representations at all—they do not bear any isomorphic relation to the serial reasoning steps we attribute to people from a folk psychological perspective. The light-duty theorist can take comfort in the fact that such networks underlie many of the most striking recent advances in artificial intelligence, including speech recognition (Hinton et al., 2012), face recognition (Parkhi, Vedaldi, & Zisserman, 2015), abstract problem solving—as deployed in games like chess and Go (Silver et al., 2016)—and pattern recognition more generally (Schmidhuber, 2015).

A second contemporary paradigm for explaining human cognition appeals to Bayesian models of probabilistic inference. Within such frameworks, different kinds of representations are hypothesized to underlie a person's knowledge in different domains (Tenenbaum et al., 2011). The point is to understand the transitions among those representations as obeying Bayesian principles of probabilistic inference. Within some Bayesian models of cognition, tree-structured representations are used; in others, two-dimensional spaces or grids are invoked, or representations resembling graphs (Tenenbaum et al., 2011, p. 1281). According to (the Bayesians) Tanenbaum et al., "Our best accounts of people's mental representations . . . resemble simpler versions of how scientists represent the same domains" (p. 1281). It is no presumption of Bayesian approaches that the representations they posit will, in general, bear transparent semantic relationships to the that-clauses of useful folk psychological talk. Again we have a flourishing research program that is not tethered to the core commitment of heavy-duty approaches. (Similar points apply to yet another popular paradigm for understanding perception and cognition: the *predictive processing theory* (Clark, 2013, 2015; Hohwy, 2013). No part of that framework assumes that the mental representations involved in such predictions correspond in any close way to our folk psychological ascriptions.)

A second avenue of response for the light-duty theorist, in the face of heavy-duty incredulity, is to grant the lack of an explanation for the dispositions we ascribe with folk psychological terms, while countering that the heavy-duty approach offers only a pseudo-explanation. Churchland & Sejnowski (1989), for instance, lampoon the heavy-duty approach to explaining human behavior by comparing it to nineteenth-century homuncular embryology, which joined the ancients in explaining the complex structure found in organisms by positing sperm which already possess the same structure in a smaller form. According to such theories, a sperm is a miniature human that, like a sponge in water, simply expands during its time in the womb (p. 161). Churchland & Sejnowski complain that Fodorian heavy-duty views explain the kind of complex linguistic behavior shown by humans—including rational inference as described via language—by appeal to mental states that have the very structure and inferential characteristics we are seeking to explain in linguistic behavior. The kinds of sentences people can say and comprehend is systematic, the Fodorian observes; so we posit a structure in the mind that is itself systematic in the very same ways. This is not unlike explaining the ten fingers and toes of adult humans by positing ten fingers and toes on a tiny human within the sperm. The heavy-duty view of human cognition is consistent with, and even "predicts," human linguistic behavior in all the ways that the homuncular theory of embryonic development predicts the growth and appearance of adult human beings. The mere fact that a *post hoc* story can be concocted that is consistent with the facts as we already knew them to be is not reason to give it special credence.

There is, of course, much more to be said on each side of the debate between heavy- and light-duty views. My aim has been to explain the nature of the debate and make room for light-duty views, without trying to settle things one way or the other. In my view, agnosticism concerning the cognitive ontology responsible for the dispositions we ascribe with folk psychological talk is reasonable at our stage of inquiry; that tilts me toward a light-duty view. But my arguments in this book won't assume either approach. My concern in the balance of this chapter is, first, to show how the project of explaining imagination differs as a function of one's being either heavy or light duty in orientation; and, second, to respond to some objections concerning the general project of explaining one folk psychological state ("imagining") in terms of a collection of others.

2.6 Explaining Imagination for Light-Duty Theorists

Supposing that one has a light-duty view of folk psychological ontology, what does it mean to explain imagination in terms of a more basic collection of folk psychological states? It means that the abilities and dispositions we attribute and

predict by ascribing imaginings to a person can *alternatively* be attributed and predicted by ascribing certain collections of beliefs, judgments, intentions, desires, and so on—all while remaining agnostic about the underlying cognitive ontology corresponding to such attributions. Consider again the folk psychological state of *suspecting*. We can, on the one hand, attribute certain dispositions to a person by saying that he *suspects that he left the stove on*. Doing so will allow us to predict and explain his behavior in various ways. On the other hand, we can attribute him the very same set of dispositions by saying that he *believes that it is somewhat likely that he left the stove on*. From the perspective of a light-duty view, neither form of ascription has greater ontological *oomph*; both attribute the same set of dispositions; both latch on to the same pattern in human behavior and inference (Dennett, 1991). For the light-duty theorist, there is no ontological dispute between the two ways of speaking—no turf battle to be waged between the notions of belief and suspicion. The phrases "Jones believes it is somewhat likely that *p*" and "Jones suspects that *p*" describe the same state of affairs. (As we will see, this is not so for the heavy-duty theorist.)

The light-duty theorist can, however, maintain that the ascription involving the word 'belief' makes use of a more general *notion*, insofar as we ascribe beliefs to people at times when it would not be appropriate to ascribe them a (mere) suspicion. By contrast, any case where a suspicion is ascribed will also be one where we could have ascribed a less than certain belief. This is the asymmetry noted in section 1.9. We can posit that there is a state of Jones in virtue of which he has those dispositions, in each case. But, for the light-duty theorist, there is no more reason to call that state "the belief that it is somewhat likely that *p*," than there is to call it "the suspicion that *p*." The light-duty theorist suspects that the notions of belief and suspicion will both have fallen out of the picture by the time we have a plausible, empirically supported theory of the state.

In many cases where two folk psychological terms serve to attribute the same dispositions and enable the same predictions, their doing so is fairly obvious. We saw this with the notions of thankfulness, regret, and suspicion, in Chapter 1. Matters are more interesting in the case of imagination. For it is *not* always easy to see how ascribing an imagining could amount to ascribing the same set of dispositions that we might with some collection of other psychological states, such as beliefs and desires. If it were, no one would raise an eyebrow at this book's core thesis. The trend in philosophy has instead been to think of imagination as a *sui generis* folk psychological state—one that, *unlike* suspecting, or being thankful, or regretting—cannot be analyzed in terms of other more general folk psychological notions such as belief, judgment, intention, and desire. To say, as a light-duty theorist, that imagination cannot be reduced to other folk psychological states is just to say that, try as we might, we cannot find a satisfying translation of the platitudes and dispositions associated with imagination-ascriptions to platitudes and dispositions we attribute with sets of other more general folk psychological

terms. It is to say that the phenomena—both mental and behavioral—we predict and explain with imagination-talk cannot alternatively be predicted and explained with belief, desire, intention, judgment, and decision-talk. Much of the work of later chapters is to show that there is in fact no such barrier; we can indeed capture the explanatory and predictive power of imagination-talk in terms of talk of beliefs, desires, and intentions (and their occurrent counterparts).

If a plausible analysis of imagination can be given along these lines, light-duty theorists should take interest. For even if explaining imagination in terms of more basic folk psychological states only amounts to showing how one set of platitudes and disposition-attributions can be translated into another, this still serves as a (surprising, to most) elucidation of imagination. Imagination is then no longer a *sui generis* mental phenomenon. A unification of one set of dispositions with another, broader set, is an explanatory unification, in Kitcher's (1981) sense (see Chapter 1). Note that the situation would be entirely different if we had no prior, independent understanding of belief, desire, intention, and so on. It would, for instance, be of far less interest to show how imagination-talk can be translated into talk of three newly invented states, described herein for the first time. The point is not simply that there is another conceivable set of states that could do the explanatory work that *sui generis* imaginings supposedly do. The key to the light-duty explanation lies in assimilating imagination-talk to talk of states we already believe in, understand, and ascribe in myriad conditions. *That* is how we reduce our stock of primitive notions.

Second, this kind of light-duty explanation has the advantage of being insulated from tumultuous debates in empirical psychology concerning the nature, format, and use of mental representations in human reasoning. Should it turn out that there is no such thing as the Belief Box or Desire Box—because there exist no mental representations with the kind of semantics and functional roles assumed by heavy-duty views—the light-duty explanation of imagination in terms of other folk psychological states retains its relevance.

Third, like anyone else, light-duty theorists expect attributions of beliefs, desires, decisions, judgments, and so on, to map, however noisily, on to *something* in the world, be it brain states, brain-body-environment pairings, or patterns of activation in neural networks—something that explains why a person has the dispositions we ascribe to him when we ascribe the state. What those things are, if their hunch is correct, just won't be all that similar to the sentences we use to ascribe folk psychological states. If cases of imagining can be understood in more basic light-duty terms, then the search for imagination's causal bases can be merged with the more general project of understanding the causal bases of the dispositions we associate with ascriptions of beliefs, desires, decisions, judgments, and so on. Questions about imagination are thereby reduced to questions about these other mental states. Here even the eliminativist about folk psychological states can take interest; for to eliminate the most basic folk psychological

states—beliefs, desires, and so on—by discovering new and better explanatory kinds will now be to eliminate imagination as well. The eliminativist will have one fewer ontological dangler.

2.6.1 Objections to this Form of Explanation, from a Light-Duty Perspective

There are worries one may nevertheless have about this sort of explanation, pitched in light-duty terms. One objection grants that the relationship mapped in Chapter 1 (Fig. 1.1b) holds, insofar as belief, desire, and intention are collectively more basic than imagination. However, it maintains that there is a sense in which imagination remains *un*reduced on such an account, precisely because (unlike suspicion), imagination is identified with a heterogeneous disjunction of different kinds of states. Arguably, where a certain type of state is identified with a heterogeneous collection of states in different token instances, the "higher level" state remains unreduced (Fodor, 1974). This is most often said to be the case when the kind-to-be-reduced enables us to make counterfactual-supporting generalizations and predictions we could not otherwise make. In such situations, the higher level kind retains an ontological significance of its own, even if, in token instances, we can perhaps do the same explanatory work by attributing some other kind of state in its stead.

Now, as it happens, I don't think that imagining *is* something like a counterfactual-supporting psychological natural kind. I think that (A-)imagining is any episode of rich, elaborated, epistemically safe thought about the possible, unreal, or fantastical. I think that's all we mean by 'imagining' when it's used in the ordinary folk psychological sense of 'imagining' captured by entries 2, 3, and 4 for 'imagine' in the *Oxford English Dictionary* (see Chapter 1). So understood, there is no reason to expect deep unity to the causal-functional profile of imaginative episodes. An imagining that occurs during a daydream can have a quite different causal role than one with the same content that occurs in the context of hypothetical reasoning, or when enjoying a fiction. (I will return to this point below.)

But even if imagining were a homogeneous counterfactual-supporting kind and, as such, retained a kind of independent ontological status, this would not stand in the way of our *explaining* imagination in terms of other psychological states. For there can be *explanatory reductions* that are not ontological reductions. Characterizing acts of imagining in other, more basic folk psychological terms provides an understanding of imagination that we previously lacked, even if one remains committed to the existence of imagination as a natural kind. The explanatory reduction allows us see how, by giving an artificial system beliefs, desires, and intentions of the right kind, we can endow it with an ability to imagine. The value of such an explanation only increases if the disjunction of states with which

imagining is characterized is not *wildly* disjunctive—if, instead, there is a smallish set of strategies for converting imagination-talk, in its paradigmatic instances, to talk of other states. That is the sort of picture I will defend by book's end.

Much the same response can be made to the objector who claims that, just as imagination-talk can be analyzed in terms of belief-, desire-, and intention-talk, so too can belief-talk be analyzed in terms of desire-, intention-, and imagination-talk. (You can take your pick of which of belief, desire, or intention gets analyzed in terms of the other two notions plus imagination. The challenge is simply that the kind of reduction proposed for imagination can be run with respect to one of the reducing states as well, with imagination serving as a primitive in *that* reduction.)[3] I don't, myself, find it at all likely that *plausible* redescriptions of this sort will be forthcoming. Which combination of desires, intentions, and imaginings will play the explanatory role of the belief that my name is 'Peter'? But let the so-motivated seek them out and convince us otherwise. If it turns out that such redescriptions are available, I would have to abandon my claim that belief, desire, and intention are collectively *more basic* than imagination. But we would still have available an explanation of imagination in other folk psychological terms. Learning that these terms are interdefinable in such ways (if they are) is to gain an important insight into the nature of the states to which they refer. And, again, appreciating the availability of such redescriptions allows us to see things—and to draw explanatory connections—that we couldn't before.

2.7 Explaining Imagination for Heavy-Duty Theorists

We've seen that when two folk psychological states have the same associated platitudes and dispositions, the light-duty view is not forced to a decision about which sort of state the person is *really* in when we ascribe one of those states. Both ascriptions point to the same place: a single set of dispositions, the causal bases of which we know not. One notion might provide explanatory leverage on the other. But, on the light-duty view, there is no deeper fact of the matter concerning which kind of state the person is in.

By contrast, the question of which attitude is ontologically *real* becomes legitimate and indeed pressing from the perspective of heavy-duty views. After all, they see psychological attitudes, taken toward concrete mental representations, as being the internal states that explain the dispositions we attribute with folk psychological talk. The heavy-duty theorist cannot lightly duplicate causes—admitting, for instance, both suspicion and belief "boxes" in the mind—in the same way light-duty theorists happily admit descriptions involving 'suspicion' and

[3] This challenge was put to me by Shen-yi Liao and Neil Van Leeuwen over lunch one day.

'belief that it is somewhat likely' as being ontologically on a par. This means that, if ascriptions of beliefs, desires, intentions (and their occurrent counterparts) really can do all the same explanatory work as ascriptions of imaginings, the heavy-duty theorist is forced to a decision on whether imagining (*qua* psychological attitude) really exists.

How will that decision be made? Consider the more neutral case of belief and suspicion: what should the heavy-duty theorist say is the *psychological attitude* that serves as the referent for ascriptions of both beliefs that it is somewhat likely that *p* and suspicions that *p*? Belief seems like the natural choice, if only because it is the more general notion. We will be able to appeal to belief in explanatory contexts including and beyond those where suspicion is an appropriate term. Why bring suspicion into our cognitive ontology, after all, if all the causal work it would do, and then some, can be done by a single psychological attitude of belief? The less than certain aspect of suspicion is accommodated through an adjustment in the content of a corresponding belief. Someone who suspects that *p*, the heavy-duty theorist can say, takes the psychological attitude of belief toward the mental sentence: it is somewhat likely that *p*. Now extend this line of thought to imagination. If the heavy-duty theorist is already committed to beliefs, desires, and intentions, and if those psychological attitudes can do all the explanatory work of imaginings *and more*, then imagination (as a psychological attitude) arrives on the chopping block.

Matters are not so straightforward, however. The nature and number of psychological attitudes is, for the heavy-duty theorist, a matter for empirical inquiry. As much as one might value parsimony in a theory (*modulo* the complexity of biological organisms), we can imagine evidence from neuropsychology that would warrant a *prima facie* less parsimonious cognitive architecture. Returning to the case of belief and suspicion, we might discover that some individuals who never show less than full certainty—political pundits, say—have a neural infarct that renders them incapable of mere suspicion. Their black-and-white views, it turns out, are a result not of careful deliberation but of dead neural tissue in Broadmann Area 10. Correlations between neural lesions at a specific site and a complete lack of suspicions might give us some reason to think that suspicion is, in fact, a distinct cognitive attitude—one that can blink out while belief chugs forward. So, while heavy-duty theorists may provisionally, on grounds of parsimony, favor views that explain both belief- and suspicion-talk in terms of a single cognitive attitude of belief, they can also leave the door open to expanding their cognitive ontologies in light of the right kind of evidence.

Imagination again presents an interesting test case, as most people haven't seen a way for cognitive attitudes like belief and desire to do the causal or explanatory work demanded by ascriptions of imaginings. If we are *already* heavy-duty theorists and cannot, from the armchair, see how more basic folk psychological terms could be used to attribute the dispositions and abilities associated with

imagination, the inference to a distinct psychological attitude of imagination—with its corresponding Imagination Box—will feel inevitable. However, we can now see that, even if it *can* be shown (for example, by me, in the balance of this book) that psychological attitudes to which heavy-duty theorists are independently committed—viz., belief, desire, and intention—are able to do the explanatory work set out for the Imagination Box, there are reasons a heavy-duty theorist might still favor a cognitive architecture that contains an Imagination Box.

The case I gave as an example, involving suspicion and political pundits, was admittedly far-fetched. However, more plausible examples have been put to work in the imagination literature. People with autism spectrum disorder (ASD) have been shown to have deficits both in their ability to engage in group pretenses and in their understanding of other minds more generally—even while maintaining high cognitive capacities in some other domains. Nichols & Stich (2003) and Currie & Ravenscroft (2002) both argue that this pattern of deficits suggests a cognitive-level dissociation between imagination and belief. They propose that their theories, which posit a distinct cognitive attitude (hereafter, a "DCA") of imagination, are better placed to explain the phenomena than accounts that posit no such distinct attitude (Nichols & Stich, 2003; Currie & Ravenscroft, 2002). (Their DCA is equivalent to my notion of a "distinct psychological attitude.") This is the sort of surprising data that *could* weigh in favor of positing a DCA of imagination (or *sui generis* imaginative states), even if, in principle, imagination-talk can be replaced with belief-, desire-, and intention-talk. However, I argue in Chapter 8 that the pattern of deficits seen in ASD offers no special support for the idea that there is a distinct cognitive attitude of imagination.

A second reason a heavy-duty theorist might posit a psychological attitude of imagination, even when belief, desire, and intention can potentially do the same explanatory and predictive work, is that the theory invoking the Imagination Box is simpler or more powerful. Of course, identifying the simplest—*qua* most time- and energy-efficient—cognitive architecture is never straightforward. To know with any certainty which proposal is more parsimonious in the relevant sense requires more than counting boxes and arrows. It requires knowing a great deal about the actual implementation of our cognitive capacities, and the costs—evolutionarily, ecologically, and metabolically—of developing and using those capacities. In many cases, weighing in on such matters with confidence will require us to know far more about the neural implementation of our mental capacities than is now understood. Arguments from parsimony are nevertheless compelling when one view attributes states to people not attributed by the other *and where both views otherwise attribute all the same states* (in terms of contents *and* attitudes). I argue in later chapters that my (imaginative-state-free) proposals are more parsimonious in this robust sense in their explanations of pretense and our engagements with fiction. It's not only the case that we can do without a *sui generis* attitude of imagination; those who posit such an attitude must, in addition

to that attitude, also posit all the same beliefs, desires, and intentions as I do. (The relevance of parsimony is more difficult to determine in situations where one view posits, say, an additional attitude not posited by the second, while the second view posits more complex contents for certain states than the first.)

Of equal importance to simplicity for a theory's power is its precision. Here more can be said on behalf of eliminating a psychological attitude of imagination. When deciding whether to include an Imagination Box in our ontology, we have to ask whether doing so enables less *noisy* predictive and explanatory generalizations than a corresponding architecture involving only belief, desire, and intention. In the case of belief and desire, we have fairly sturdy *ceteris paribus* theorems that allow us to predict and explain behavior, such as: if someone desires that *p* and believes that φ-ing will make it the case that *p*, then she will endeavor to φ, provided she has no stronger countervailing desires. There are exceptions to this sort of generalization—hence the *ceteris paribus*. People have seizures, trip over roots, or are simply too drunk or too tired to φ. These phenomena constitute noise in the pattern picked out by the theorem; yet all sides tend to agree that *ceteris paribus* generalizations remain genuinely explanatory, as they appear in all but the most basic sciences (Dennett, 1991; Fodor, 1987).

We can suppose that there are also theorems, or *ceteris paribus* predictive patterns, involving the term 'imagines.' Those patterns will also be subject to exceptions; they will be noisy. Should it turn out that they are *much more* noisy than patterns we can exploit when redescribing the same behavior in other folk psychological terms, then the replacing terms (i.e., those of the redescription) will have greater explanatory power—they will predict more things correctly, more of the time. Imagination will have been explained in the sense that the considerable noise within explanations involving 'imagines' will have been reduced. From a heavy-duty perspective, such reductions in noise are reasons to think that the psychological attitudes posited by the noise-reducing theory better match reality. From a light-duty perspective, less noisy explanatory patterns are epistemically preferable in allowing for more predictive success, provided they are not much more difficult to exploit (Dennett, 1991).

There is reason to think that the folk psychological theorems that invoke 'imagining' are indeed noisy and subject to exceptions, relative to those involving terms like 'belief' and 'desire.' It is common to encounter proposals about what a person who imagines that *p* *can* or *may* do; but it is rare to find claims about what they *will* or *must do*, *ceteris paribus*. For example: it has been said that a person who imagines that *p* and believes that if *p* then *q* will tend to imagine that *q*; this platitude finds its way into formal characterizations of imagination's role in hypothetical reasoning (Carruthers, 2006; Nichols & Stich, 2000; Van Leeuwen, 2014, p. 795). And yet: we may at any time imagine that *p* and believe that if p then q, without then imagining that *q*—and not because we had a seizure, tripped on a root, were distracted, or too drunk. We may fail to imagine that *q* simply because

doing so doesn't fit our goals or interests at the time, or because it never occurred to us to be so realistic in what we were imagining. For instance, I believe that if I arrive to teach unprepared, the class session will be tedious and stressful. And yet, I just imagined that I arrived at class unprepared and had a great, lively discussion. Have I flouted a norm? Hardly. Unrealistic imagining is par for the course; it's part of what makes imagining *imagining*. To say that common sense provides *ceteris paribus* generalizations about the causal or inferential role of imaginings is an overstatement.

Kathleen Stock makes much the same point, noting that:

> There are barely any platitudes about the causal role of the imagination, implicit in ordinary language . . . unlike other mental entities such as belief and desire, the functional role of imagining is relatively unclear . . . There is little distinctive behavior associated with either imaginings with particular contents, or imaginings generally . . . Equally, there seem to be few predictable generalizations connecting imagining to other mental states or events. (2017, p. 4)

Note the difference between Stock's plausible claim there are no platitudes about the *causal role* of imagination, and the false claim that there are no platitudes *whatsoever* about imagination. While there are indeed plenty of platitudes about imagination—many of which were reviewed in Chapter 1—Stock's point is that such platitudes don't coalesce to paint a clear picture of the causal-functional role of an imaginative state. In support she lists several examples where an imaginative state with the content *p* has a causal role in one context quite unlike what it has in another. Her conclusion is that conceptual analysis—which limits itself to facts about a state's causal role known by any competent speaker—will be of limited use in analyzing imagination (2017, p. 5).

Here is a different conclusion we might reasonably draw: we do not, with imagination, have our hands on a single psychological kind (at least, not from a causal-functional point of view),[4] but instead a heterogeneous assortment of different, more basic folk psychological states which *do* have comparatively clean causal roles. Supposing *this* hypothesis is true, our ability to replace imagination-talk with talk of these other kinds of states will greatly improve our predictive and explanatory abilities. This point holds relevance for both heavy- and light-duty views.

[4] Imaginings do retain a *kind* of unity on my view, relative to the two senses of the word "imagine" distinguished in Chapter 1. They are all cases of rich, elaborated, epistemically safe thought about the merely possible, fantastical, and so on (in the case of A-imaginings); and they are all cases of seemingly image-like thought (in the case of I-imaginings). The presence of this kind of unity, in each case, explains why we are tempted to analyze imaginings as a class in the first place. Yet neither is a kind of unity that enables much in the way of behavioral and inferential predictions and explanations, as the characterizations do not suggest a single causal-functional role for either type of imagining.

An alternative, more common, reaction to the kind of heterogeneity Stock observes is to propose a "cleaning up" of the notion of imagination—keying one's predictive generalizations involving 'imagines' to situations where imagination operates in a (supposed) *default mode* (Williamson, 2016, p. 116),[5] or when particular *constraints are applied to it* (Kind, 2016a). That is fair game, but it increases the complexity of one's overall picture. The generalizations and predictive heuristics we employ no longer simply involve the term 'imagines' in its ordinary (not-cleaned-up) sense. So the relevant patterns and generalizations no longer fall naturally out of the platitudes competent speakers will accept about imagination. Making use of them will require explicitly articulating and empirically validating a new, more complex vocabulary for describing human inferential and behavioral dispositions. The resulting theory will no longer have folk psychology's simplicity and implicit validation-through-use on its side. By all means, if new constructs are indeed needed to explain the phenomena, we should get to work in formulating and testing them. But if we can, with equal or better predictive and explanatory success, employ an existing folk psychological vocabulary we already successfully use in other contexts . . . well then that's much better!

Finally, there remains an easily-overlooked challenge worth noting to any heavy-duty account positing a distinct cognitive attitude of imagination. As we saw earlier, the platitudes surrounding imagination do not paint a clear or univocal picture of its causal role. Yet when heavy-duty theorists turn the cognitive attitude of imagination into an explanatory *posit*, they have to give it a fairly precise causal role: it must have the role of causing whatever it is they have called on it to explain. In all likelihood, some of the messiness in the pre-reflective, folk psychological notion of imagination will have been trimmed off. What are we to do with the clippings? We can't sweep them into the trash without a second thought. Cleaning up the concept of imagination, so as to give it a respectable causal role, does not make the shorn behaviors and mental phenomena *disappear*. If, for instance, the psychological attitude of imagination does a great job in explaining highly constrained hypothetical reasoning while leaving fantastical daydreams a mystery, then the psychological attitude of imagination doesn't explain all of what we want explained by a theory of imagination. These gaps must be acknowledged when the theory is compared to others that *do* explain the full set of phenomena—perhaps by finding imagination to be a heterogeneous kind, constituted by a collection of more basic folk psychological states.

[5] "Left to itself, the imagination develops the scenario in a reality-oriented way, by default" (Williamson, 2016, p. 116).

2.8 Summary

When we think of folk psychological states in light-duty terms, we see them as sets of mental and behavioral dispositions whose causal bases we know not. Working from the philosopher's armchair, we needn't be bashful about our knowledge of folk psychological states *so conceived*. We want to be discussing imagination itself, not just the concepts surrounding imagination, after all. The light-duty view shows how this can be done, without our proposing to limn the structure of the mind in the process. If it turns out that the generalizations and patterns associated with imagination-talk are a mess, and even self-contradictory, we have good reason to seek other ways of attributing the same dispositions and capacities with better-behaved, more basic folk psychological terms. But even if we think imaginings are a well-behaved folk psychological kind, we can still arrive at an explanation of imagination by seeing how behaviors and cognitive capacities associated with imagination-talk can be alternatively described and cataloged through the use of other familiar mental state terms. Such an explanation is all the more powerful if the patterns and generalizations invoked are less noisy and have greater predictive precision than those featuring 'imagines.'

The light-duty view's conservativeness about mental ontology also facilitates a kind of explanatory pluralism. Our ability to articulate questions about imagination's relation to other states in light-duty terms allows us to pitch present debates in a relatively theory-neutral way. Cognitive boxologies can be rejected as wrongheaded by one party, for instance, while the question of imagination's relation to—and possible reducibility to—other folk psychological states remains a shared theoretical question.

The heavy-duty approach, by contrast, pronounces on the contents of certain mental representations realized in the brain—namely, those that explain our having of the dispositions ascribed by folk psychological talk—and the psychological attitudes taken toward them. And, at times, it appeals to surprising empirical results in support of doing so. Yet the heavy-duty theorist can also argue against imagination's reducibility to (or explainability in terms of) other folk psychological states in just the same way as the light-duty theorist. If it turns out that we cannot capture the behavioral patterns and dispositions associated with imagination-talk in more basic folk psychological terms, the light-duty theorist concludes that imagination is a *sui generis* folk psychological mental phenomenon; the heavy-duty theorist concurs but goes further in holding that we have defeasible evidence for the existence of a distinct psychological attitude of imagination—an "Imagination Box."

On the other hand, if the arguments in later chapters succeed, then we can indeed replace imagination-talk with talk of beliefs, desires, and intentions. This will be reason to think there is no such psychological attitude (or DCA) as the

heavy-duty theorist proposes. Imagination would not thereby be *eliminated*. Rather, a particular theoretical construct that sometimes goes by the name of 'imagination'—one that only occurs with certain heavy-duty ontologies—would be eliminated. Imagination, as a folk psychological phenomenon, would persist. On a light-duty ontology, nothing at all gets eliminated if my arguments in later chapters succeed. Instead, we come to see that the notions of belief, desire, and intention can serve in explanations of what it is to imagine.